

IMS  
Version 11

## *System Administration*





IMS  
Version 11

## *System Administration*



**Note**

Before using this information and the product that it supports, be sure to read the general information under “Notices” on page 847.

This edition applies to IMS Version 11 (program number 5635-A02) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1974, 2015.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

## About this information. . . . . xiii

|                                                         |      |
|---------------------------------------------------------|------|
| Prerequisite knowledge . . . . .                        | xiii |
| IMS function names used in this information . . . . .   | xiii |
| How new and changed information is identified . . . . . | xiv  |
| Accessibility features for IMS Version 11 . . . . .     | xiv  |
| How to send your comments . . . . .                     | xv   |

## | Changes to the IMS library . . . . . xvii

---

## Part 1. Introduction to IMS system administration . . . . . 1

### Chapter 1. Introduction to the IMS system . . . . . 3

|                                                 |    |
|-------------------------------------------------|----|
| Overview of administrative activities . . . . . | 3  |
| IMS environments and configurations . . . . .   | 5  |
| DB/DC environment . . . . .                     | 6  |
| DBCTL environment . . . . .                     | 11 |
| DCCTL environment . . . . .                     | 13 |
| DB batch environment. . . . .                   | 17 |
| TM batch environment . . . . .                  | 18 |
| IMSpdex overview . . . . .                      | 19 |

### Chapter 2. Concepts for the system administrator . . . . . 31

|                                                                   |    |
|-------------------------------------------------------------------|----|
| System support for application programs . . . . .                 | 31 |
| Dynamic allocation with IMS . . . . .                             | 32 |
| Base Primitive Environment overview . . . . .                     | 33 |
| Tracing BPE components . . . . .                                  | 35 |
| Common Service Layer overview . . . . .                           | 38 |
| CSL in an IMSpdex . . . . .                                       | 39 |
| CSL managers . . . . .                                            | 41 |
| A single point of control (SPOC) program in CSL . . . . .         | 45 |
| REXX SPOC API. . . . .                                            | 48 |
| Global online change . . . . .                                    | 49 |
| ACB library member online change . . . . .                        | 50 |
| Global TM resource management . . . . .                           | 51 |
| Automatic RECON loss notification overview . . . . .              | 51 |
| Configuring an IMSpdex with CSL. . . . .                          | 52 |
| CSL configuration examples . . . . .                              | 54 |
| Common Queue Server overview . . . . .                            | 62 |
| Queue structures . . . . .                                        | 65 |
| Resource structures. . . . .                                      | 66 |
| CQS structure functions . . . . .                                 | 66 |
| CQS recovery functions . . . . .                                  | 67 |
| Type-2 command environment . . . . .                              | 68 |
| Extended Terminal Option . . . . .                                | 68 |
| APPC . . . . .                                                    | 69 |
| Security for dependent region processing . . . . .                | 69 |
| MPP scheduling in DB/DC and DCCTL environments. . . . .           | 70 |
| Scheduling application programs against unavailable data. . . . . | 71 |
| Fast Path . . . . .                                               | 72 |

|                                                                        |    |
|------------------------------------------------------------------------|----|
| Fast Path in a DBCTL environment . . . . .                             | 73 |
| Automated operator application programs . . . . .                      | 73 |
| System logging and processing continuity . . . . .                     | 74 |
| Checkpointing . . . . .                                                | 75 |
| Locking mechanisms and database integrity in DB/DC and DCCTL . . . . . | 76 |
| Exit points and routines . . . . .                                     | 77 |
| Data Capture exit routines . . . . .                                   | 77 |
| HALDB Partition Selection exit routines . . . . .                      | 78 |
| The z/OS Automatic Restart Manager (ARM) . . . . .                     | 79 |

### Chapter 3. Documenting the IMS system. . . . . 81

|                                                      |    |
|------------------------------------------------------|----|
| Extracting requirements for the IMS system . . . . . | 81 |
| Participating in design reviews . . . . .            | 82 |
| Establishing naming conventions . . . . .            | 83 |
| Using a data dictionary . . . . .                    | 84 |
| Documenting the system characteristics . . . . .     | 85 |

### Chapter 4. Introduction to IMS operations and recovery . . . . . 87

|                                                 |     |
|-------------------------------------------------|-----|
| Understanding the operations task . . . . .     | 87  |
| Automated operations. . . . .                   | 88  |
| Understanding the recovery task . . . . .       | 89  |
| Example: a system without recovery . . . . .    | 89  |
| Requirements for a recoverable system . . . . . | 90  |
| Steps in the process of recovery . . . . .      | 90  |
| The mechanisms of recovery in IMS . . . . .     | 91  |
| Backward recovery . . . . .                     | 115 |
| Forward recovery . . . . .                      | 116 |
| Shutdown of the IMS system . . . . .            | 116 |
| Recovery during restart . . . . .               | 117 |
| Complications in recovery . . . . .             | 117 |
| What IMS cannot recover . . . . .               | 118 |
| Overhead of recovery. . . . .                   | 118 |
| IMS Application Menu . . . . .                  | 119 |

---

## Part 2. IMS system administration considerations and tasks . . . . . 123

### Chapter 5. z/OS interface considerations. . . . . 125

|                                                                                     |     |
|-------------------------------------------------------------------------------------|-----|
| z/OS interface considerations for IMS . . . . .                                     | 125 |
| Preventing installation problems . . . . .                                          | 125 |
| Setting up JCL . . . . .                                                            | 125 |
| Configuring JES for dependent regions. . . . .                                      | 126 |
| Keeping some required nonstandard z/OS macros in their original libraries . . . . . | 126 |
| Updating the IBM-supplied z/OS Program Properties Table . . . . .                   | 126 |
| Installing required IMS links to z/OS . . . . .                                     | 129 |
| Installing the type 2 SVC module. . . . .                                           | 131 |
| Binding the channel-to-channel (CTC) channel-end appendage . . . . .                | 132 |

|  |                                                                                              |     |
|--|----------------------------------------------------------------------------------------------|-----|
|  | Installing the resource cleanup module . . . . .                                             | 132 |
|  | IMS abend formatting module. . . . .                                                         | 134 |
|  | Adding the offline dump formatting routine to<br>the print dump exit control table . . . . . | 136 |
|  | Binding the DBRC type 4 SVC . . . . .                                                        | 136 |
|  | Authorizing IMS system data sets in the<br>Authorized Program Facility . . . . .             | 136 |
|  | Updating the APPC/MVS administration dialog                                                  | 137 |
|  | RRS Archive Log Stream considerations . . . . .                                              | 137 |
|  | System Management Facility DDCONS<br>parameter considerations . . . . .                      | 137 |
|  | z/OS interface considerations for IRLM . . . . .                                             | 137 |

## **Chapter 6. VTAM interface considerations . . . . . 139**

Setting the Network Control Program (NCP) delay 140

## **Chapter 7. CSL administration . . . . . 141**

|  |                                                                    |     |
|--|--------------------------------------------------------------------|-----|
|  | Using the z/OS Automatic Restart Manager with<br>the CSL . . . . . | 141 |
|  | Administering global online change . . . . .                       | 142 |
|  | Enabling global online change . . . . .                            | 142 |
|  | Disabling global online change . . . . .                           | 142 |
|  | Maintaining an IMSplex with mixed online<br>change scope . . . . . | 143 |
|  | Administering automatic RECON loss notification                    | 143 |
|  | Monitoring an IMSplex . . . . .                                    | 144 |
|  | Diagnosing IMSplex problems. . . . .                               | 145 |

## **Chapter 8. CSL ODBM administration 147**

|  |                                                               |     |
|--|---------------------------------------------------------------|-----|
|  | ODBM client registration in a CSL . . . . .                   | 147 |
|  | ODBM and RRS . . . . .                                        | 148 |
|  | ODBM and ODBA . . . . .                                       | 148 |
|  | Configuring ODBA application servers to use<br>ODBM . . . . . | 149 |
|  | ODBM message routing . . . . .                                | 149 |
|  | ODBM and security . . . . .                                   | 150 |

## **Chapter 9. CSL OM administration 151**

|  |                                                                             |     |
|--|-----------------------------------------------------------------------------|-----|
|  | CSL OM command routing. . . . .                                             | 151 |
|  | Issuing commands through the OM API . . . . .                               | 152 |
|  | CSL OM audit trail . . . . .                                                | 153 |
|  | Configuring the logstream for the CSL OM<br>audit trail . . . . .           | 154 |
|  | OM audit log record formats . . . . .                                       | 155 |
|  | Printing OM log records. . . . .                                            | 155 |
|  | CSL OM command security . . . . .                                           | 157 |
|  | RACF authorization for commands entered<br>through OM. . . . .              | 158 |
|  | IMS commands, RACF access authorities and<br>resource names table . . . . . | 159 |
|  | CSL OM Security user exit routine . . . . .                                 | 168 |

## **Chapter 10. CSL RM administration 169**

|  |                                                                                         |     |
|--|-----------------------------------------------------------------------------------------|-----|
|  | Information managed by the CSL RM . . . . .                                             | 169 |
|  | Maintaining global information for databases,<br>DEDB areas, and transactions . . . . . | 170 |
|  | Maintaining message destination resource<br>information . . . . .                       | 172 |

|  |                                                                                             |     |
|--|---------------------------------------------------------------------------------------------|-----|
|  | Maintaining global information for the<br>OLCSTAT data set. . . . .                         | 173 |
|  | Maintaining IMSplex-wide parameters . . . . .                                               | 173 |
|  | How IMS maintains global information for<br>sysplex serialized program management . . . . . | 173 |
|  | Resource structure duplexing requirements for CSL<br>RM. . . . .                            | 173 |
|  | How the CSL RM repopulates a resource structure                                             | 173 |
|  | How z/OS rebuilds a resource structure . . . . .                                            | 174 |

## **Chapter 11. CSL SCI administration 175**

|  |                            |     |
|--|----------------------------|-----|
|  | CSL SCI security . . . . . | 175 |
|--|----------------------------|-----|

## **Chapter 12. CQS administration . . . 177**

|  |                                                                                  |     |
|--|----------------------------------------------------------------------------------|-----|
|  | Recording information necessary for starting CQS                                 | 177 |
|  | Establishing client connection to CQS during failed<br>client takeover . . . . . | 177 |
|  | Authorizing access to CQS . . . . .                                              | 178 |
|  | Authorizing CQS registration . . . . .                                           | 178 |
|  | Authorizing connections to CQS structures . . . . .                              | 178 |
|  | Using structure alter for CQS . . . . .                                          | 179 |
|  | Using CQS system checkpoint. . . . .                                             | 180 |
|  | Using CQS structure checkpoint . . . . .                                         | 181 |
|  | Preventing CQS structure full . . . . .                                          | 182 |
|  | CQS structure overflow function . . . . .                                        | 182 |
|  | CQS structure full monitoring . . . . .                                          | 183 |
|  | Using structure full monitoring with CQS<br>structure overflow. . . . .          | 184 |
|  | Rebuilding structures in CQS . . . . .                                           | 185 |
|  | z/OS system-managed rebuild and CQS . . . . .                                    | 185 |
|  | CQS-managed rebuild . . . . .                                                    | 185 |
|  | Initiating structure rebuild with z/OS and CQS                                   | 186 |
|  | CQS structure repopulation. . . . .                                              | 186 |
|  | CQS structure recovery . . . . .                                                 | 187 |
|  | CQS structure copy . . . . .                                                     | 188 |
|  | z/OS structure duplexing for CQS . . . . .                                       | 188 |

## **Chapter 13. IMS service considerations . . . . . 191**

|  |                                                                                                         |     |
|--|---------------------------------------------------------------------------------------------------------|-----|
|  | Types of service SMP/E SYSMODs . . . . .                                                                | 191 |
|  | Service SYSMOD packaging . . . . .                                                                      | 191 |
|  | Obtaining IMS service . . . . .                                                                         | 192 |
|  | Maintenance recommendations . . . . .                                                                   | 193 |
|  | Assessing your readiness to apply maintenance                                                           | 193 |
|  | General maintenance recommendations for a<br>production system. . . . .                                 | 194 |
|  | General maintenance recommendations for IMS<br>service distributed after testing cycle begins . . . . . | 194 |
|  | Sample implementation of maintenance<br>recommendations . . . . .                                       | 194 |
|  | Installing IMS service on a single system . . . . .                                                     | 195 |
|  | RECEIVE/APPLY/ACCEPT (standard sequence)                                                                | 195 |
|  | ACCEPT without APPLY (pregeneration mode)                                                               | 196 |
|  | ACCEPT before APPLY (for service that requires<br>a system definition) . . . . .                        | 199 |
|  | Installing IMS service in an IMSplex. . . . .                                                           | 200 |
|  | Common installation and maintenance issues. . . . .                                                     | 201 |
|  | Preventing regression of SYSMODs in<br>APPLY-only status by an IMS system definition . . . . .          | 201 |

|                                                                                    |            |                                                                                 |     |
|------------------------------------------------------------------------------------|------------|---------------------------------------------------------------------------------|-----|
| Generating JCL to build non-system definition target libraries . . . . .           | 202        | Using IRLM with database-level sharing . . . . .                                | 238 |
| Applying maintenance for the IVP dialog . . . . .                                  | 202        | Block-level sharing . . . . .                                                   | 239 |
| Upgrading z/OS . . . . .                                                           | 202        | Multiple IRLMs in sysplex data sharing . . . . .                                | 239 |
| Ensuring proper SYSLIB concatenation . . . . .                                     | 202        | Setting up IRLM procedures . . . . .                                            | 240 |
| Interpreting binder return codes properly . . . . .                                | 203        | System initialization for IRLM. . . . .                                         | 242 |
| Assembling and binding a sample exit routine using SMP/E . . . . .                 | 204        | Defining an IRLM for sysplex data sharing . . . . .                             | 243 |
| Migrating to a new version of IMS . . . . .                                        | 204        | Data sharing at the database level with update activity . . . . .               | 243 |
| <b>Chapter 14. Planning for shared queues . . . . .</b>                            | <b>205</b> | Data sharing at the database level with multiple readers . . . . .              | 244 |
| The role of CQS in a shared-queues environment                                     | 205        | Data sharing at the block level . . . . .                                       | 245 |
| The IMS role in a shared-queues environment                                        | 205        | Tailoring IMS systems that share data . . . . .                                 | 246 |
| How IMS registers interest in queues . . . . .                                     | 206        | Installation tasks . . . . .                                                    | 246 |
| Message flow in a shared-queues environment                                        | 208        | DBRC and IRLM support for batch systems in a data-sharing environment . . . . . | 247 |
| Units of work (UOW) tracking . . . . .                                             | 208        | Excluding a database from data sharing . . . . .                                | 248 |
| Terminal autologon . . . . .                                                       | 209        | Tailoring execution JCL . . . . .                                               | 248 |
| Message and program-to-program switches . . . . .                                  | 209        | Tailoring the z/OS operating system . . . . .                                   | 248 |
| Dynamic control blocks . . . . .                                                   | 209        | Coordinating VSAM data set definitions with share options . . . . .             | 249 |
| IMS Queue Manager . . . . .                                                        | 209        | Starting and stopping data sharing . . . . .                                    | 249 |
| MTO messages . . . . .                                                             | 210        | Starting the IRLM . . . . .                                                     | 249 |
| Serial application processing . . . . .                                            | 210        | Stopping the IRLM . . . . .                                                     | 250 |
| Serial transaction processing . . . . .                                            | 210        | Starting and stopping database processing. . . . .                              | 250 |
| Conversational transactions in a shared-queues environment. . . . .                | 211        | Monitoring data-sharing systems . . . . .                                       | 252 |
| Configuring a shared-queues environment. . . . .                                   | 211        | Obtaining the status of IRLM activity . . . . .                                 | 252 |
| Enabling shared queues . . . . .                                                   | 214        | Displaying components and resources . . . . .                                   | 252 |
| Defining the CFRM policy . . . . .                                                 | 215        | Monitoring information in the RECON data set . . . . .                          | 254 |
| Defining a z/OS log stream . . . . .                                               | 217        | Monitoring structures on a coupling facility . . . . .                          | 254 |
| Defining CQS parameters . . . . .                                                  | 218        | Using DBRC to control database allocation and access . . . . .                  | 257 |
| IMS parameters for shared queues . . . . .                                         | 218        | Online DBRC commands . . . . .                                                  | 258 |
| Using the Common Queue Server . . . . .                                            | 219        | Improving database availability . . . . .                                       | 259 |
| Monitoring and requeuing structures using the Queue Control Facility . . . . .     | 222        | Controlling online change . . . . .                                             | 259 |
| Accessing CQS with IMS commands . . . . .                                          | 223        | Reorganizing databases . . . . .                                                | 259 |
| List structures . . . . .                                                          | 224        | Making an online database image copy. . . . .                                   | 260 |
| Using associated printers . . . . .                                                | 226        | Maintaining normal operations . . . . .                                         | 260 |
| Planning for IMS front-end switch . . . . .                                        | 227        | Scheduling a batch updater. . . . .                                             | 260 |
| Determining eligibility for sysplex-wide processing                                | 227        | Reinstating an online updater . . . . .                                         | 261 |
| Managing APPC and OTMA messages in a sysplex environment. . . . .                  | 228        | Transferring update capability. . . . .                                         | 262 |
| Managing the IMS dead-letter queue . . . . .                                       | 229        | Reorganizing a database. . . . .                                                | 262 |
| Using the Expedited Message Handler . . . . .                                      | 229        | Making an image copy with exclusive control . . . . .                           | 263 |
| Requeuing suspended transactions . . . . .                                         | 230        | Making an image copy of a shared database . . . . .                             | 263 |
| Scheduling AOI transactions . . . . .                                              | 230        | Lowering the share level of a database . . . . .                                | 264 |
| Planning for MSC in a shared-queues environment                                    | 230        | Recovering from failures . . . . .                                              | 265 |
| Planning for RSR in a shared-queues environment                                    | 231        | Managing system logs . . . . .                                                  | 265 |
| Tuning for performance in a shared-queues environment. . . . .                     | 232        | Planning recovery procedures . . . . .                                          | 266 |
| Recommendation for Fast Path transactions in a shared-queues environment . . . . . | 232        | Recovery facilities . . . . .                                                   | 266 |
| <b>Chapter 15. Data sharing in IMS environments . . . . .</b>                      | <b>233</b> | Recovery without DBRC. . . . .                                                  | 268 |
| Data sharing overview . . . . .                                                    | 233        | Restart after IMS failure . . . . .                                             | 268 |
| DBRC and data sharing support . . . . .                                            | 234        | Restart after DBRC failure . . . . .                                            | 269 |
| Database integrity without data sharing support                                    | 236        | Recovery involving IRLM . . . . .                                               | 269 |
| How applications access data . . . . .                                             | 237        | Data sharing in a sysplex environment . . . . .                                 | 271 |
| How IMS systems share databases . . . . .                                          | 237        | Sysplex data-sharing concepts and terminology                                   | 271 |
|                                                                                    |            | XRF and sysplex data sharing . . . . .                                          | 274 |
|                                                                                    |            | Sample sysplex data-sharing configurations . . . . .                            | 274 |
|                                                                                    |            | When to use sysplex data sharing . . . . .                                      | 279 |
|                                                                                    |            | Modifying batch jobs for sysplex data sharing                                   | 279 |
|                                                                                    |            | Calculating the size of coupling facility structures. . . . .                   | 279 |



|                                                         |     |
|---------------------------------------------------------|-----|
| Changing the size of OSAM and VSAM structures . . . . . | 283 |
| Recovering from sysplex data-sharing failures . . . . . | 283 |

## Chapter 16. Planning for VTAM generic resource groups . . . . . 289

|                                                                         |     |
|-------------------------------------------------------------------------|-----|
| Requirements for using VTAM generic resource groups . . . . .           | 289 |
| VTAM generic resource group restrictions . . . . .                      | 289 |
| VTAM generic resource affinity . . . . .                                | 289 |
| Creating a VTAM generic resource group . . . . .                        | 291 |
| Specifying APPC generic resource names . . . . .                        | 291 |
| Initiating sessions with IMS in a VTAM generic resource group . . . . . | 291 |
| XRF and VTAM generic resources . . . . .                                | 292 |
| Changing the logon procedure . . . . .                                  | 292 |
| Overriding the APPLID field . . . . .                                   | 292 |
| Determining when affinities are terminated . . . . .                    | 293 |
| Terminating affinities that persist . . . . .                           | 294 |
| Dropping an IMS system from a generic resource group . . . . .          | 294 |
| Resetting terminal status . . . . .                                     | 294 |
| Logging on after IMS failure with affinity . . . . .                    | 295 |
| Controlling VTAM generic resource member selection . . . . .            | 295 |
| Ensuring consistency across the IMSplex . . . . .                       | 296 |
| Bypassing affinity management during the IMS ESTAE process . . . . .    | 297 |

## Chapter 17. Planning for Transaction Manager resources in an IMSplex . . . 299

|                                             |     |
|---------------------------------------------|-----|
| Resource name uniqueness . . . . .          | 299 |
| Resource type consistency . . . . .         | 300 |
| Global callable services . . . . .          | 300 |
| Transaction Manager resources . . . . .     | 300 |
| TM resources: APPC descriptors . . . . .    | 300 |
| TM resources: VTAM LTERMs. . . . .          | 301 |
| TM resources: MSNAMEs . . . . .             | 301 |
| TM resources: VTAM terminal nodes . . . . . | 302 |
| TM resources: transactions . . . . .        | 302 |
| TM resources: user names . . . . .          | 303 |
| TM resources: user IDs . . . . .            | 304 |

|                                                                   |     |
|-------------------------------------------------------------------|-----|
| How RM and the resource structure impact IMS activities . . . . . | 304 |
|-------------------------------------------------------------------|-----|

## Chapter 18. IMS security . . . . . 307

|                                                                      |     |
|----------------------------------------------------------------------|-----|
| Overview of DB/DC and DCCTL security. . . . .                        | 307 |
| DB/DC and DCCTL resources that can be protected . . . . .            | 307 |
| Defining security during DB/DC and DCCTL system definition . . . . . | 308 |
| Security facilities for DB/DC and DCCTL resources . . . . .          | 308 |
| Designing security for IMS DB/DC and DCCTL . . . . .                 | 310 |
| Limiting access from a terminal . . . . .                            | 310 |
| Authorizing transactions and commands . . . . .                      | 312 |
| Using RACF to protect physical terminals . . . . .                   | 314 |
| Security considerations for the master terminal . . . . .            | 315 |
| Security for AO application programs . . . . .                       | 315 |
| Security for Time-Controlled Operations . . . . .                    | 320 |

|                                                                                |     |
|--------------------------------------------------------------------------------|-----|
| Security for Fast Path application programs . . . . .                          | 321 |
| Security for JMP application programs . . . . .                                | 321 |
| Security and CPI-C driven application programs . . . . .                       | 321 |
| Security for ODBA application programs . . . . .                               | 322 |
| Use of the RACF data space . . . . .                                           | 322 |
| Security in MSC and shared-queues environments . . . . .                       | 322 |
| Security for APPC/IMS . . . . .                                                | 326 |
| Security for ETO terminals . . . . .                                           | 327 |
| Securing DB/DC and DCCTL dependent regions and their resources . . . . .       | 327 |
| Security for OTMA . . . . .                                                    | 330 |
| Security for IMS Connect . . . . .                                             | 330 |
| Activating IMS security for DB/DC and DCCTL environments . . . . .             | 331 |
| Defining the SECURITY macro for DB/DC and DCCTL environments . . . . .         | 331 |
| Preparing security exit routines . . . . .                                     | 331 |
| Preparing a RACF security plan . . . . .                                       | 332 |
| Enabling and disabling APSB SAF security . . . . .                             | 335 |
| Controlling security during system startup for DB/DC and DCCTL . . . . .       | 335 |
| Implementing security changes online in DB/DC and DCCTL environments . . . . . | 337 |
| Controlling security violations in DB/DC and DCCTL environments . . . . .      | 337 |
| Considering other access control methods . . . . .                             | 338 |
| Physical security . . . . .                                                    | 338 |
| Using display bypass and password masking in DB/DC and DCCTL . . . . .         | 338 |
| Protecting your resources . . . . .                                            | 339 |
| Encryption: an alternative to access control . . . . .                         | 341 |
| Security considerations for a DBCTL environment . . . . .                      | 341 |
| DBCTL resources that can be protected . . . . .                                | 341 |
| DBCTL security choices made during system definition. . . . .                  | 342 |
| Security facilities and security types for DBCTL resources . . . . .           | 342 |
| Design considerations for DBCTL security. . . . .                              | 343 |
| Activating IMS DBCTL security . . . . .                                        | 344 |
| Controlling system startup for DBCTL security . . . . .                        | 346 |
| Implementing DBCTL security changes online . . . . .                           | 346 |
| Controlling DBCTL security violations . . . . .                                | 347 |

## Chapter 19. Controlling IMS . . . . . 349

|                                                              |     |
|--------------------------------------------------------------|-----|
| Monitoring the system . . . . .                              | 349 |
| Processing IMS system log information for analysis . . . . . | 349 |
| Using IMS system log utilities. . . . .                      | 349 |
| IMS database productivity tools . . . . .                    | 350 |
| Recovering the system . . . . .                              | 350 |
| Modifying and controlling system resources . . . . .         | 351 |
| Controlling data sharing. . . . .                            | 352 |
| Controlling data sharing using DBRC . . . . .                | 352 |
| Monitoring the data-sharing system . . . . .                 | 352 |
| Controlling log data set characteristics . . . . .           | 353 |
| Connecting and disconnecting subsystems. . . . .             | 353 |



## Chapter 20. Starting or restarting IMS 355

## Chapter 21. Shutting down IMS. . . . 357

## Chapter 22. Testing the system. . . . 359

|                                                                                                     |     |
|-----------------------------------------------------------------------------------------------------|-----|
| The need for a test system . . . . .                                                                | 359 |
| Setting up a test system . . . . .                                                                  | 360 |
| Setting up a test database . . . . .                                                                | 361 |
| Testing operational procedures . . . . .                                                            | 361 |
| Monitoring in IMS test environments . . . . .                                                       | 361 |
| Monitoring in a DB/DC environment . . . . .                                                         | 362 |
| Ensuring network readiness . . . . .                                                                | 362 |
| Network testing in DB/DC and DCCTL environments . . . . .                                           | 363 |
| Testing in a DBCTL environment . . . . .                                                            | 363 |
| Testing in a DCCTL environment . . . . .                                                            | 363 |
| IMS testing aids . . . . .                                                                          | 364 |
| Simulating online execution with Batch Terminal Simulator in DB/DC and DCCTL environments . . . . . | 364 |
| Online testing of MFS formats in DB/DC and DCCTL environments . . . . .                             | 364 |
| Using dynamic resource definition and online change for testing . . . . .                           | 365 |
| Program testing using SYSIN/SYSOUT in DB/DC and DCCTL environments . . . . .                        | 366 |
| Network testing using Teleprocessing Network Simulator in DB/DC and DCCTL environments . . . . .    | 367 |
| Performance and stress testing using the Queue Control Facility . . . . .                           | 367 |

## Chapter 23. Collecting and interpreting IMS monitoring data . . . . 369

|                                                                                   |     |
|-----------------------------------------------------------------------------------|-----|
| Establishing monitoring procedures . . . . .                                      | 369 |
| Establishing performance objectives . . . . .                                     | 370 |
| Planning for workload management . . . . .                                        | 371 |
| Deciding on monitoring activities and techniques . . . . .                        | 375 |
| Tools for detailed monitoring . . . . .                                           | 378 |
| IMS Monitor . . . . .                                                             | 378 |
| z/OS Generalized Trace Facility (GTF) . . . . .                                   | 380 |
| z/OS Component Trace (CTRACE) Service . . . . .                                   | 380 |
| Obtaining program isolation and lock traces . . . . .                             | 381 |
| IMS Trace facility . . . . .                                                      | 382 |
| Analyzing performance in OTMA message routing . . . . .                           | 383 |
| Coordinating performance information in an MSC network . . . . .                  | 383 |
| Monitoring Fast Path systems in DB/DC and DCCTL environments . . . . .            | 384 |
| Transaction flow in DB/DC and DCCTL environments . . . . .                        | 384 |
| The IMS Monitor in DB/DC and DCCTL environments . . . . .                         | 388 |
| Monitoring transaction-level statistics . . . . .                                 | 388 |
| Monitoring procedures in a DBCTL environment . . . . .                            | 389 |
| Establishing performance objectives in a DBCTL environment . . . . .              | 390 |
| Deciding on monitoring activities and techniques in a DBCTL environment . . . . . | 392 |

## Chapter 24. Modifying the system design . . . . . 395

|                                                                                |     |
|--------------------------------------------------------------------------------|-----|
| Assessing application changes . . . . .                                        | 395 |
| Introducing changed applications in an active IMS system . . . . .             | 398 |
| Resizing the inactive ACB library online . . . . .                             | 398 |
| Migrating static allocation of the ACB library to dynamic allocation . . . . . | 399 |
| Planning for system definition changes . . . . .                               | 399 |
| Controlling system definition processing . . . . .                             | 399 |
| Determining the type of system definition required . . . . .                   | 400 |
| Changing the network definition in DB/DC and DCCTL environments . . . . .      | 400 |
| Coordinating system definition and current security definitions . . . . .      | 401 |
| Making system tuning changes . . . . .                                         | 401 |
| Resource utilization changes . . . . .                                         | 401 |
| Application and database design changes . . . . .                              | 402 |
| Communication design changes in DB/DC and DCCTL environments . . . . .         | 402 |
| Managing online system definition changes . . . . .                            | 402 |
| Deciding if system modifications can use online change . . . . .               | 402 |
| Planning considerations for online change . . . . .                            | 403 |
| Performing capacity planning . . . . .                                         | 405 |

## Chapter 25. Tuning the system . . . . 407

|                                                                                         |     |
|-----------------------------------------------------------------------------------------|-----|
| Managing performance . . . . .                                                          | 408 |
| Change management . . . . .                                                             | 409 |
| Design variables . . . . .                                                              | 409 |
| Types of prediction . . . . .                                                           | 410 |
| Initializing z/OS and IMS parameters for tuning . . . . .                               | 410 |
| Assigning z/OS dispatching priorities . . . . .                                         | 410 |
| Choosing IMS options for performance . . . . .                                          | 412 |
| Avoiding contention for IMS resources (excluding buffer pools) . . . . .                | 414 |
| IMS buffer pool tuning . . . . .                                                        | 416 |
| Trade-offs between I/O controlled by IMS and paging . . . . .                           | 420 |
| Minimizing path length . . . . .                                                        | 423 |
| Considerations for the communications network in DB/DC and DCCTL environments . . . . . | 423 |
| Guidelines for IMS system data set placement . . . . .                                  | 424 |
| I/O subsystem configuration . . . . .                                                   | 425 |
| Application optimization in DB/DC and DCCTL environments . . . . .                      | 425 |
| DL/I considerations . . . . .                                                           | 425 |
| Planning for performance in a shared-queues environment . . . . .                       | 426 |
| Identifying and correcting performance problems . . . . .                               | 426 |
| Examining paging rates . . . . .                                                        | 427 |
| Detecting processor resource problems . . . . .                                         | 428 |
| Tuning to remove I/O resource contention . . . . .                                      | 428 |
| Communication subsystem contention in DB/DC and DCCTL environments . . . . .            | 428 |
| IMS message processing in DB/DC and DCCTL environments . . . . .                        | 430 |
| Input queuing and scheduling/termination in DB/DC and DCCTL environments . . . . .      | 431 |

|                                                                                       |            |                                                                               |            |
|---------------------------------------------------------------------------------------|------------|-------------------------------------------------------------------------------|------------|
| Program load and initialization . . . . .                                             | 433        | Controlling print data sets . . . . .                                         | 471        |
| Program execution times . . . . .                                                     | 434        | Express alternate PCBs . . . . .                                              | 472        |
|                                                                                       |            | XRF environments. . . . .                                                     | 472        |
|                                                                                       |            | Understanding allocation errors . . . . .                                     | 472        |
| <b>Chapter 26. Making online changes</b>                                              | <b>437</b> |                                                                               |            |
| Overview of the local online change function. . . . .                                 | 438        | <b>Chapter 28. Accessing external systems from within IMS</b>                 | <b>475</b> |
| IMS environments that support the online change function . . . . .                    | 440        |                                                                               |            |
| Commands applicable to online change . . . . .                                        | 441        | <b>Part 3. DBRC administration.</b>                                           | <b>477</b> |
| Changing the system definition online . . . . .                                       | 441        |                                                                               |            |
| Changing programs and transactions . . . . .                                          | 442        | <b>Chapter 29. Overview of DBRC</b>                                           | <b>479</b> |
| Implementing security changes online . . . . .                                        | 442        | DBRC components . . . . .                                                     | 480        |
| Changing MFS formats . . . . .                                                        | 442        | DBRC tasks . . . . .                                                          | 481        |
| Making online changes for DEDBs . . . . .                                             | 443        | DBRC commands and API requests . . . . .                                      | 481        |
| Making online changes for HALDBs. . . . .                                             | 443        |                                                                               |            |
| Making local online changes in an XRF complex                                         | 443        | <b>Chapter 30. Recording and controlling log information</b>                  | <b>485</b> |
| Online change performance considerations . . . . .                                    | 444        | Changing log records in the RECON data sets . . . . .                         | 486        |
| Overview of the global online change function . . . . .                               | 444        | Archiving an OLDS . . . . .                                                   | 486        |
| Commands applicable to global online change                                           | 444        | DBRC log-related commands . . . . .                                           | 488        |
| Changing the system definition online . . . . .                                       | 445        |                                                                               |            |
| Changing programs and transactions . . . . .                                          | 446        | <b>Chapter 31. Recovering databases</b>                                       | <b>489</b> |
| Implementing security changes online . . . . .                                        | 446        |                                                                               |            |
| Changing MFS formats . . . . .                                                        | 446        | <b>Chapter 32. How DBRC helps in recovery</b>                                 | <b>493</b> |
| Making online changes for DEDBs . . . . .                                             | 447        | Generating recovery JCL . . . . .                                             | 493        |
| Making online changes for HALDBs. . . . .                                             | 447        | Validating utility JCL. . . . .                                               | 494        |
| Changing runtime resource definitions . . . . .                                       | 447        |                                                                               |            |
| Making local online changes in an XRF complex                                         | 448        | <b>Chapter 33. Recording information about opening and updating databases</b> | <b>497</b> |
| Making global online changes in a sysplex that uses XRF . . . . .                     | 449        |                                                                               |            |
| Global online change after XRF takeover and DBCTL standby emergency restart . . . . . | 449        | <b>Chapter 34. Supporting data sharing</b>                                    | <b>499</b> |
| Online change performance considerations . . . . .                                    | 450        | Levels of data sharing . . . . .                                              | 499        |
| Making local online changes in a sysplex . . . . .                                    | 450        | Recording data-sharing information in the RECON data sets . . . . .           | 499        |
| Enabling the global online change function . . . . .                                  | 451        | Assigning a sharing level with DBRC . . . . .                                 | 499        |
| Initiating a global online change . . . . .                                           | 452        |                                                                               |            |
| Changing or adding IMS.ACBLIB members online . . . . .                                | 454        | <b>Chapter 35. DBRC support for Remote Site Recovery</b>                      | <b>501</b> |
| Displaying status or attribute information about IMS systems in the IMSplex . . . . . | 456        |                                                                               |            |
| Terminating an online change . . . . .                                                | 456        | <b>Chapter 36. Supporting IMSplexes</b>                                       | <b>503</b> |
| Handling errors for global online change commands . . . . .                           | 457        |                                                                               |            |
| Global online change scenarios . . . . .                                              | 457        | <b>Chapter 37. Defining DBRC to IMS</b>                                       | <b>505</b> |
| Performing a cold start . . . . .                                                     | 463        |                                                                               |            |
| Determining when to perform a cold start. . . . .                                     | 463        | <b>Chapter 38. Registering databases and database data sets</b>               | <b>507</b> |
| Cleaning up after a failed restart attempt . . . . .                                  | 464        |                                                                               |            |
|                                                                                       |            | <b>Chapter 39. Planning for recovery</b>                                      | <b>509</b> |
| <b>Chapter 27. Printing output with IMS Spool API</b>                                 | <b>465</b> | Setting up recovery mechanisms . . . . .                                      | 509        |
| Design and operational considerations . . . . .                                       | 465        | Recovery facilities . . . . .                                                 | 509        |
| The IMS Spool API as a data manager . . . . .                                         | 467        | Recovery without DBRC. . . . .                                                | 511        |
| Print data set characteristics . . . . .                                              | 468        | Restart after an IMS failure. . . . .                                         | 512        |
| The Change (CHNG) call . . . . .                                                      | 468        | Restart after a DBRC failure . . . . .                                        | 512        |
| The Set Options (SETO) call . . . . .                                                 | 469        | Recovery involving IRLM configurations . . . . .                              | 512        |
| The Output DD statement . . . . .                                                     | 469        | Batch backout . . . . .                                                       | 512        |
| Writing data to the IMS Spool API . . . . .                                           | 470        |                                                                               |            |
| The Insert call . . . . .                                                             | 470        |                                                                               |            |
| The PURG call . . . . .                                                               | 470        |                                                                               |            |
| ROLL and ROLB calls . . . . .                                                         | 471        |                                                                               |            |
| SETS, SETU, and ROLS calls . . . . .                                                  | 471        |                                                                               |            |
| Special considerations—descriptors allowed . . . . .                                  | 471        |                                                                               |            |

## **Chapter 40. Considerations for a DBRC system . . . . . 513**

|                                                                                 |     |
|---------------------------------------------------------------------------------|-----|
| General considerations for using DBRC. . . . .                                  | 513 |
| Data set naming conventions . . . . .                                           | 513 |
| Naming convention for image copy data sets                                      | 514 |
| Naming convention for duplicate image copy data sets . . . . .                  | 515 |
| Naming convention for change accumulation data sets . . . . .                   | 515 |
| Database backup copies . . . . .                                                | 515 |
| Image copy utilities (DFSUDMP0, DFSUDMT0, DFSUICP0) . . . . .                   | 516 |
| Concurrent image copy . . . . .                                                 | 519 |
| Creating image copy data sets for future use and reuse. . . . .                 | 520 |
| Controlling the number of image copies managed. . . . .                         | 520 |
| Recovery period of image copy data sets and GENMAX . . . . .                    | 521 |
| Reusing image copy data sets . . . . .                                          | 522 |
| HSSP image copy . . . . .                                                       | 523 |
| HISAM copies (DFSURUL0 and DFSURRL0) . . . . .                                  | 523 |
| Nonstandard image copy data sets . . . . .                                      | 524 |
| Frequency and retention of backup copies. . . . .                               | 524 |
| Log record change accumulation . . . . .                                        | 525 |
| Condensing the accumulated SLDS or RLDS (change accumulation) . . . . .         | 525 |
| When is change accumulation required? . . . . .                                 | 526 |
| DBDS group considerations . . . . .                                             | 529 |
| DB groups . . . . .                                                             | 530 |
| DBRC groups . . . . .                                                           | 530 |
| DBRC considerations for HALDB Online reorganization . . . . .                   | 531 |
| RSR tracking system . . . . .                                                   | 533 |
| Maintaining recovery-related records affected by online reorganization. . . . . | 533 |
| Online reorganization process . . . . .                                         | 533 |

## **Chapter 41. Initializing and maintaining the RECON data sets . . . 535**

|                                                               |     |
|---------------------------------------------------------------|-----|
| Planning considerations for the RECON data sets               | 535 |
| Initializing the RECON data sets for the first time . . . . . | 536 |
| Avoiding RECON data set contention problems                   | 536 |
| Initially accessing the RECON data sets . . . . .             | 544 |
| Records in the RECON data sets . . . . .                      | 545 |
| RECON header records . . . . .                                | 546 |
| Log data set records . . . . .                                | 546 |
| Database recovery records . . . . .                           | 548 |
| Mapping the records in the RECON data set . . . . .           | 555 |
| Maintaining the RECON data sets . . . . .                     | 556 |
| Backing up the RECON data sets. . . . .                       | 556 |
| Deleting unnecessary RECON records . . . . .                  | 556 |
| Reorganizing the RECON data sets . . . . .                    | 559 |
| Replacing damaged RECON data sets . . . . .                   | 560 |
| Recovering the RECON data sets. . . . .                       | 560 |
| Replacing a discarded RECON data set. . . . .                 | 561 |
| RECON loss notification. . . . .                              | 562 |
| DBRC security . . . . .                                       | 563 |
| Tracking changes made to the RECON data sets                  | 566 |

## **Chapter 42. Hints and tips for DBRC 567**

|                                                                                      |     |
|--------------------------------------------------------------------------------------|-----|
| Changing the RECON data sets to output time stamps in local time of origin . . . . . | 567 |
| Locating the last SLDS stop time in the RECON data set . . . . .                     | 567 |
| Adjusting GENMAX when it is reached or it is too high . . . . .                      | 569 |
| Closing an open online PRILOG record. . . . .                                        | 570 |
| Working with subsystem records (SSYS) . . . . .                                      | 572 |
| Removing authorization inconsistency between the SSYS from DB/AREA records . . . . . | 572 |
| Restarting the change accumulation process for logs . . . . .                        | 573 |
| Restarting the change accumulation work when it states nothing to process . . . . .  | 573 |
| Moving log data sets . . . . .                                                       | 574 |
| Cataloging data sets . . . . .                                                       | 574 |
| Performing multiple cold starts in a test environment. . . . .                       | 575 |
| Reducing potential RECON data set enqueue problems . . . . .                         | 576 |

## **Chapter 43. Resource names for command authorization . . . . . 579**

## **Part 4. IMS system recovery. . . . 583**

## **Chapter 44. Extended Recovery Facility Overview . . . . . 585**

|                                                           |     |
|-----------------------------------------------------------|-----|
| Overview of XRF . . . . .                                 | 585 |
| XRF in an IMSplex . . . . .                               | 585 |
| Installation types that benefit from XRF . . . . .        | 586 |
| XRF concepts and terminology . . . . .                    | 586 |
| XRF complexes. . . . .                                    | 588 |
| DBCTL capabilities . . . . .                              | 591 |
| XRF takeover . . . . .                                    | 591 |
| XRF requirements . . . . .                                | 595 |
| Component roles in the XRF process . . . . .              | 597 |
| Contribution of IMS to the XRF process . . . . .          | 597 |
| Contributions of z/OS and the z/OS elements               | 600 |
| Establishing surveillance for XRF. . . . .                | 606 |
| Phases of the XRF process . . . . .                       | 608 |
| Initialization phase of the XRF process . . . . .         | 608 |
| Synchronization phase of the XRF process. . . . .         | 610 |
| Tracking phase of the XRF process . . . . .               | 611 |
| Takeover phase of the XRF process . . . . .               | 612 |
| Post-takeover phase of the XRF process . . . . .          | 621 |
| Termination phase of the XRF process . . . . .            | 623 |
| Cycle of XRF phases . . . . .                             | 624 |
| Organization of XRF complexes . . . . .                   | 624 |
| One XRF complex with one CPC . . . . .                    | 624 |
| One XRF complex with two CPCs . . . . .                   | 626 |
| One XRF complex with two CPCs and a non-XRF IMS . . . . . | 626 |
| Two XRF complexes with three CPCs . . . . .               | 627 |
| Two XRF complexes with four CPCs . . . . .                | 629 |
| Planning an XRF complex . . . . .                         | 631 |
| Limitations of XRF . . . . .                              | 632 |
| Non-XRF workload on the alternate IMS . . . . .           | 632 |

|                                                                 |     |
|-----------------------------------------------------------------|-----|
| Using the Intersystem Communication link . . . . .              | 632 |
| Terminals in an XRF complex . . . . .                           | 633 |
| How takeover affects a terminal user . . . . .                  | 638 |
| How VTAM ownership affects terminal switching. . . . .          | 642 |
| Planning for VTAM ownership of class-1 terminals . . . . .      | 642 |
| NCP planning considerations for XRF with USERVAR . . . . .      | 643 |
| Performance considerations. . . . .                             | 643 |
| z/OS planning considerations . . . . .                          | 644 |
| Preparing the system for XRF . . . . .                          | 645 |
| Tailoring the IMS execution JCL . . . . .                       | 646 |
| Coding IMS system definition macro statements for XRF . . . . . | 646 |
| Defining IMS.PROCLIB members for XRF . . . . .                  | 649 |
| Defining the VTAM USERVAR table. . . . .                        | 656 |
| Placement of IMS data sets in the XRF configuration . . . . .   | 656 |

## Chapter 45. Remote Site Recovery

### overview . . . . . 663

|                                                                              |     |
|------------------------------------------------------------------------------|-----|
| RSR overview . . . . .                                                       | 663 |
| Requirements for using RSR . . . . .                                         | 665 |
| Basic components of RSR . . . . .                                            | 665 |
| RSR processing. . . . .                                                      | 672 |
| Determining the extent of recovery . . . . .                                 | 672 |
| Defining an RSR environment . . . . .                                        | 674 |
| XRF and RSR . . . . .                                                        | 680 |
| Defining an RSR environment with XRF . . . . .                               | 681 |
| Data sharing and RSR . . . . .                                               | 682 |
| DRD and RSR . . . . .                                                        | 683 |
| Tracking an IMSplex . . . . .                                                | 685 |
| RSR log management. . . . .                                                  | 685 |
| Example of an RSR complex . . . . .                                          | 687 |
| Coordinated IMS and DB2 for z/OS recovery support . . . . .                  | 688 |
| Installing RSR . . . . .                                                     | 691 |
| Hardware replication. . . . .                                                | 691 |
| Software replication . . . . .                                               | 691 |
| Running IMS workload on multiple z/OS images in an RSR environment . . . . . | 695 |
| Initializing RSR. . . . .                                                    | 696 |
| Initializing the active site . . . . .                                       | 696 |
| Initializing the remote site . . . . .                                       | 697 |
| IMS error handling for RSR for the active site . . . . .                     | 698 |
| IMS error handling for RSR for the remote site . . . . .                     | 700 |
| Establishing IMS security for RSR . . . . .                                  | 703 |

## Chapter 46. Recovery in an IMSplex 705

## Part 5. Using IMS reports . . . . . 707

### Chapter 47. DB Monitor reports . . . . . 709

|                                                 |     |
|-------------------------------------------------|-----|
| VSAM-Buffer-Pool report . . . . .               | 709 |
| Using the VSAM-Buffer-Pool report . . . . .     | 709 |
| Fields in the VSAM-Buffer-Pool report . . . . . | 709 |
| VSAM-Statistics report . . . . .                | 715 |
| Using the VSAM-Statistics report. . . . .       | 715 |

|                                                                  |     |
|------------------------------------------------------------------|-----|
| Fields in the VSAM-Statistics report . . . . .                   | 715 |
| Database-Buffer-Pool report . . . . .                            | 720 |
| Using the Database-Buffer-Pool report . . . . .                  | 720 |
| Fields in the Database-Buffer-Pool report . . . . .              | 720 |
| Program-I/O report . . . . .                                     | 722 |
| Using the Program-I/O report. . . . .                            | 722 |
| Fields in the Program-I/O report. . . . .                        | 722 |
| DL/I-Call-Summary report . . . . .                               | 724 |
| Using the DL/I-Call-Summary report . . . . .                     | 724 |
| Fields in the DL/I-Call-Summary report . . . . .                 | 725 |
| Distribution-Appendix report . . . . .                           | 727 |
| How to generate the Distribution-Appendix report . . . . .       | 729 |
| Events that can be distributed and their default ranges. . . . . | 729 |
| Monitor-Overhead report . . . . .                                | 730 |
| Fields in the Monitor-Overhead report . . . . .                  | 730 |

### Chapter 48. IMS Monitor reports . . . . . 733

|                                                                |     |
|----------------------------------------------------------------|-----|
| Transaction flow and IMS Monitor events . . . . .              | 733 |
| IMS Monitor trace event intervals . . . . .                    | 735 |
| Overview of IMS Monitor reports . . . . .                      | 736 |
| Documenting the monitoring run. . . . .                        | 737 |
| Verifying IMS Monitor report occurrences . . . . .             | 739 |
| Monitoring activity in dependent regions . . . . .             | 740 |
| Detecting database processing intent conflicts                 | 743 |
| Examining the effects of checkpoints . . . . .                 | 743 |
| Measuring region occupancy . . . . .                           | 743 |
| Monitoring application program elapsed time . . . . .          | 744 |
| Monitoring I/O for application program DL/I Calls . . . . .    | 746 |
| Monitoring MFS activity. . . . .                               | 750 |
| Monitoring message queue handling . . . . .                    | 750 |
| Detecting checkpoint effects . . . . .                         | 751 |
| Transaction Queuing report. . . . .                            | 752 |
| Monitoring database buffers . . . . .                          | 753 |
| Monitoring line activity . . . . .                             | 754 |
| Monitoring message handling efficiency . . . . .               | 755 |
| IMS internal resource usage . . . . .                          | 755 |
| Using frequency distributions from IMS Monitor output. . . . . | 758 |
| How to get a frequency distribution output . . . . .           | 760 |
| How frequency distribution ranges are defined                  | 760 |
| Interpreting the Distribution Appendix. . . . .                | 762 |
| IMS Monitor reports for MSC . . . . .                          | 763 |
| MSC Traffic report. . . . .                                    | 764 |
| MSC Summaries report . . . . .                                 | 765 |
| MSC Queuing Summary report . . . . .                           | 766 |

### Chapter 49. IMS Monitor reports for DBCTL . . . . . 767

|                                                             |     |
|-------------------------------------------------------------|-----|
| IMS Monitor trace event intervals . . . . .                 | 768 |
| Overview of IMS Monitor reports . . . . .                   | 769 |
| Documenting the monitoring run. . . . .                     | 770 |
| Monitoring activity in dependent regions . . . . .          | 772 |
| Monitoring application program elapsed time . . . . .       | 776 |
| Monitoring I/O for application program DL/I calls . . . . . | 778 |
| Transaction Queuing report. . . . .                         | 779 |
| Monitoring database buffers . . . . .                       | 780 |

|                                                                |     |
|----------------------------------------------------------------|-----|
| IMS internal resource usage . . . . .                          | 781 |
| Using frequency distributions from IMS Monitor output. . . . . | 782 |
| Interpreting Distribution Appendix output . . . . .            | 786 |

## **Chapter 50. IMS Monitor reports for DCCTL . . . . . 787**

|                                                                |     |
|----------------------------------------------------------------|-----|
| IMS Monitor trace event intervals . . . . .                    | 787 |
| Overview of IMS Monitor reports . . . . .                      | 788 |
| Documenting the monitoring run. . . . .                        | 789 |
| Monitoring activity in dependent regions . . . . .             | 791 |
| Examining the effects of checkpoints . . . . .                 | 794 |
| Measuring region occupancy . . . . .                           | 795 |
| Monitoring application program elapsed time . . . . .          | 795 |
| Monitoring I/O for application program DL/I calls . . . . .    | 797 |
| Monitoring MFS activity. . . . .                               | 800 |
| Monitoring message queue handling . . . . .                    | 801 |
| Detecting checkpoint effects . . . . .                         | 801 |
| Transaction Queuing report. . . . .                            | 802 |
| Monitoring line activity . . . . .                             | 803 |
| Monitoring message handling efficiency . . . . .               | 804 |
| IMS internal resource usage . . . . .                          | 804 |
| Using frequency distributions from IMS Monitor output. . . . . | 805 |
| Interpreting Distribution Appendix output . . . . .            | 809 |

## **Chapter 51. //DFSSTAT reports . . . . . 811**

|                                                            |     |
|------------------------------------------------------------|-----|
| PST-Accounting report . . . . .                            | 811 |
| VSAM-Buffer-Pool report (for batch regions only) . . . . . | 812 |
| OSAM-Buffer-Pool report (for batch regions only) . . . . . | 813 |

|                                               |     |
|-----------------------------------------------|-----|
| Sequential-Buffering-Summary report . . . . . | 814 |
| Sequential Buffering Detail report . . . . .  | 816 |

## **Chapter 52. Statistical-analysis, log-transaction reports, and analyzing log records . . . . . 825**

|                                                      |     |
|------------------------------------------------------|-----|
| Statistical analysis utility reports . . . . .       | 825 |
| Calculating transaction loads . . . . .              | 826 |
| Assessing program-to-program traffic . . . . .       | 827 |
| Obtaining counts of unsent messages . . . . .        | 827 |
| Auditing critical transactions . . . . .             | 828 |
| Log transaction analysis utility reports . . . . .   | 828 |
| Examining scheduling activity. . . . .               | 829 |
| IMS accounting information . . . . .                 | 832 |
| Analyzing log records using KBLA utilities . . . . . | 833 |
| Overview of KBLA routines . . . . .                  | 833 |
| ISPF interface for KBLA . . . . .                    | 838 |

## **Part 6. Appendixes . . . . . 845**

### **Notices . . . . . 847**

|                                             |     |
|---------------------------------------------|-----|
| Programming interface information . . . . . | 849 |
| Privacy policy considerations . . . . .     | 850 |

### **Bibliography. . . . . 851**

### **Index . . . . . 853**





---

## About this information

These topics provide guidance information for managing the administration and operations of a single IMS™ or one or more IMS systems that work as a unit (an IMSplex). The topics describe designing, documenting, operating, maintaining, and recovering an IMS system, as well as the Database Recovery Control (DBRC) facility, Remote Site Recovery (RSR), and the Extended Recovery Facility (XRF). The topics also include information about the IMS Base Primitive Environment (BPE), IMS Common Queue Server (CQS), and IMS Common Service Layer (CSL), all of which can be part of an IMSplex, as well as task information for sharing data and message queues in an IMSplex.

This information is available in IBM® Knowledge Center at [www.ibm.com/support/knowledgecenter](http://www.ibm.com/support/knowledgecenter).

---

## Prerequisite knowledge

Before using this book, you should understand basic z/OS® and IMS concepts and your installation's IMS system. IMS can run in the following environments: DB Batch, DCCTL, TM Batch, DB/DC, DBCTL. You should understand the environments that apply to your installation. You should also have a basic understanding of database processing and the access methods used by DL/I. It is helpful to know the purpose of the different types of DL/I calls, the IMS application program structure, and the tasks associated with application program design.

You can learn more about z/OS by visiting the “z/OS basic skills education” topics in IBM Knowledge Center.

You can gain an understanding of basic IMS concepts by reading *An Introduction to IMS*, an IBM Press publication. An excerpt from this publication is available in the IBM Knowledge Center.

IBM offers a wide variety of classroom and self-study courses to help you learn IMS. For a complete list of courses available, go to the IMS home page at [ibm.com/ims](http://ibm.com/ims) and link to the Training and Certification page.

---

## IMS function names used in this information

In this information, the term HALDB Online Reorganization refers to the integrated HALDB Online Reorganization function that is part of IMS Version 11, unless otherwise indicated.

IMS provides an integrated IMS Connect function, which offers a functional replacement for the IMS Connect tool (program number 5655-K52). In this information, the term IMS Connect refers to the integrated IMS Connect function that is part of IMS Version 11, unless otherwise indicated.



---

## How new and changed information is identified

New and changed information in most IMS library PDF publications is denoted by a character (revision marker) in the left margin. The first edition (-00) of *Release Planning*, as well as the *Program Directory* and *Licensed Program Specifications*, do not include revision markers.

Revision markers follow these general conventions:

- Only technical changes are marked; style and grammatical changes are not marked.
- If part of an element, such as a paragraph, syntax diagram, list item, task step, or figure is changed, the entire element is marked with revision markers, even though only part of the element might have changed.
- If a topic is changed by more than 50%, the entire topic is marked with revision markers (so it might seem to be a new topic, even though it is not).

Revision markers do not necessarily indicate all the changes made to the information because deleted text and graphics cannot be marked with revision markers.

---

## Accessibility features for IMS Version 11

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use information technology products successfully.

### Accessibility features

The following list includes the major accessibility features in z/OS products, including IMS Version 11. These features support:

- Keyboard-only operation.
- Interfaces that are commonly used by screen readers and screen magnifiers.
- Customization of display attributes such as color, contrast, and font size.

### Keyboard navigation

You can access IMS Version 11 ISPF panel functions by using a keyboard or keyboard shortcut keys.

For information about navigating the IMS Version 11 ISPF panels using TSO/E or ISPF, refer to the *z/OS TSO/E Primer*, the *z/OS TSO/E User's Guide*, and the *z/OS ISPF User's Guide Volume 1*. These guides describe how to navigate each interface, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

### Related accessibility information

Online documentation for IMS Version 11 is available in IBM Knowledge Center.

### IBM and accessibility

See the *IBM Human Ability and Accessibility Center* at [www.ibm.com/able](http://www.ibm.com/able) for more information about the commitment that IBM has to accessibility.

---

## How to send your comments

Your feedback is important in helping us provide the most accurate and highest quality information. If you have any comments about this or any other IMS information, you can take one of the following actions:

- Click the **Feedback** link at the bottom of any IBM Knowledge Center topic.
- Send an email to [imspubs@us.ibm.com](mailto:imspubs@us.ibm.com). Be sure to include the book title and the publication number.



---

## Changes to the IMS library

Changes to the IMS library for IMS Version 11 include minor title changes to all the publications, merging four sets of publications, a change to how *IMS Messages and Codes* are delivered, and a reorganization of *IMS Version 11 Exit Routines*. The library is no longer provided in BookManager format. In addition, as of the May, 2012, a master index is no longer provided.

The titles of all PDF publications in the IMS Version 11 library have been changed so that the publication titles match the contents view (navigation tree) titles in IBM Knowledge Center. For example, the *IMS Diagnosis Guide* has been renamed to *IMS Diagnosis*. In addition, the titles for the three *IMS Command Reference* PDF volumes have been renamed so that their titles better reflect the contents of each publication, as follows:

- *IMS Version 11 Commands, Volume 1: IMS Commands A-M*
- *IMS Version 11 Commands, Volume 2: IMS Commands N-V*
- *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*

Four sets of publications have been streamlined:

- To eliminate redundancy and group relevant application programming planning information with associated application programming tasks, the *IMS Application Programming Planning Guide* has been merged with the *IMS Application Programming Guide*, in a new title: *IMS Version 11 Application Programming*.
- To group diagnosis reference information with associated diagnosis tasks, the *IMS Diagnosis Reference* has been merged with the *IMS Diagnosis Guide*, in a new title: *IMS Version 11 Diagnosis*.
- To group all system administration tasks, regardless of whether the tasks relate to using a single IMS or an IMSplex, the *IMS IMSplex Administration Guide* has been merged with the *IMS System Administration Guide*, in a new title: *IMS Version 11 System Administration*.
- To group system definition reference information with associated system definition tasks, the *IMS System Definition Reference* has been merged with the *IMS System Definition Guide*, in a new title: *IMS Version 11 System Definition*.

*IMS Version 11 Exit Routines* is now organized by the call interface and available services for each group of exit routines, as follows:

- IMS system exits, database manager exits, and transaction manager exit routines are grouped under “IMS control region exit routines.” These exit routines use IMS interfaces and callable services, as described in the topic “Overview of exit routines for IMS.”
- Exit routines that run in an address space with BPE support, including those for CQS, BPE-based DBRC, and most CSL components, are grouped under “Base Primitive Environment-based exit routines.” BPE-based exit routines use the interfaces and services provided by the Base Primitive Environment, as described in the topic “BPE user-supplied exit routine interfaces and services.”
- The CQS exit routines are now divided into two groups: those that run in the CQS address space and have access to BPE services are grouped under “BPE-based CQS user-supplied exit routines,” and those that execute in a CQS client address space are grouped under “CQS client exit routines.”

- Standard DBRC exit routines are grouped under “IMS system exit routines,” while the variants that run in BPE are group under “BPE-based DBRC exit routines.”
- CSL SCI exit routines are now divided into two categories: those that run in the CSL SCI address space are grouped under “BPE-based SCL SCI user exit routines,” and those that run in the address space of a registered IMSplex member are grouped under “CSL SCI IMSplex member exit routines.”
- Exit routines that have specific guidelines, including those for IMS Connect, CQS clients, and CSL SCI IMSplex members, are described in the sections for those routines.

---

## **Part 1. Introduction to IMS system administration**

This topic introduces the IMS system. It also introduces the concepts that are central to administering an IMS system.





---

# Chapter 1. Introduction to the IMS system

This topic outlines the administrative activities for an IMS system. It also introduces the concepts that are central to administering an IMS system.

---

## Overview of administrative activities

To administer an online IMS system, you must design an IMS online system, establish operating procedures that meet application requirements, maintain a production system that is responsive to end users, and integrate new applications or major design changes into the current system.

To meet these responsibilities, you must coordinate many activities that occur during an application development cycle. Performance of these activities results in:

- Documentation of the IMS network
- Specifications for IMS system definition and execution control parameters
- Procedural controls for operation
- Strategies for monitoring and audit trails

After entry into production mode, your ongoing activities support:

- Auditing operations and end user service
- Monitoring and gathering production statistics
- Establishing procedures to control changes to the online system design
- Testing the online system after application and IMS changes

The following figure is an overview of the administration activities. The activities in the first column of the figure take place during the design phase, the second column's activities occur during development of application code, and those in the third column occur during test. The vertical center line marks the transition to production mode.

Before the start of production mode, administration activities lead to two major activities:

- Generating the system and preparing JCL (job control language)
- Developing operations procedures

The items to the right of the center line in the figure reflect the ongoing system administration activities. The column headed by PRODUCTION emphasizes awareness of the day-to-day operation and performance of the system. The next column's activities are concerned with maintenance. Minor application design changes, problem resolution, and any IMS maintenance are included in this category. The column on the far right identifies the activities needed to integrate additional applications or implement an application package. For a major addition, activities are similar to those on the left-hand side of the figure. Other design changes that do not involve terminal or network modifications can be handled without shutting down the IMS. For these simple design changes, the associated activities lead to a revision of operating procedures in maintenance mode.

After startup, revisions of IMS system definition and operating procedures become key activities. In this context, you might decide to change applications during

online operation. The interpretation of system performance then becomes an important activity, supported by monitoring and performance analysis.

Each row in the figure represents a set of activities, each having its own characteristics:

- Analyzing user requirements involves examining the application documentation for the IMS function required and the expected workload.
- Collecting online requirements concerns specifications for IMS system definition, system data set allocation, and initial JCL.
- Preparing the IMS network involves interaction with system and network generation activities.
- Establishing security procedures encompasses the design and implementation of a security strategy.
- Developing an operations plan produces operations control documents and provides for audit of the control of production cycles.
- Forming a monitoring strategy results in monitoring the system and gathering performance data.
- Establishing criteria for performance leads to performance analysis and tuning activities.




Figure 1. Chart of IMS system administration activities

**Related reading:** This publication does not address the detailed planning needed to establish operating procedures. For information on establishing operating procedures, see *IMS Version 11 Operations and Automation*.

## IMS environments and configurations

IMS contains two major components that can be used together or separately: the IMS Database Manager (DB) and the IMS Transaction Manager (TM). When you use them together, they constitute the DB/DC environment.

Using the Database Manager alone, you can generate the batch environment and the database control (DBCTL) environment. Using the Transaction Manager alone, you can generate the data communication control (DCCTL) environment. Data sharing and the Extended Recovery Facility (XRF) are often considered environments, but they are special cases of the three environments listed here.

In addition to the three environments (DBCTL, DB/DC, and DCCTL), you can configure IMS as an *IMSplex*.

Use the major components of IMS to configure your environment as part of your system definition process, based on your business needs.

The DB/DC, DCCTL, and DBCTL environments are all considered *online* IMS systems.

Each of the IMS environments is a distinct combination of hardware and programs that supports distinct processing goals. The online environments and the goals they support are shown in the following table.

*Table 1. IMS online environments*

| Environments                               | Data processing goals                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DBCTL (See “DBCTL environment” on page 11) | <ul style="list-style-type: none"><li>• Process network transactions without the Transaction Manager—that is, use the Database Manager with a transaction management subsystem (for example, CICS®).</li><li>• Run batch application programs using DB batch at certain intervals (for example, process a payroll or produce an inventory report).</li><li>• Run database utilities using DB batch.</li></ul>                                                                                                                                                                                                                                                                                                                    |
| DB/DC (See “DB/DC environment”)            | <ul style="list-style-type: none"><li>• Enable terminal users to retrieve data and modify the database with satisfactory real-time performance. (Some typical applications are banking, airline reservations, and sales orders.)</li><li>• Ensure that retrieved data is current.</li><li>• Distribute transaction processing among multiple CPUs in a communications network.</li><li>• Run batch application programs using DB batch at certain intervals (for example, process a payroll or produce an inventory report).</li><li>• Run database utilities using DB batch.</li><li>• Run application programs that access external subsystems or access data in external subsystems, such as data in DB2® for z/OS.</li></ul> |
| DCCTL (See “DCCTL environment” on page 13) | <ul style="list-style-type: none"><li>• Process network transactions without the Database Manager by using the Transaction Manager with an external database management subsystem.</li><li>• Maintain system log information for restart by using DBRC.</li><li>• Run application programs that access external subsystems or access data in external subsystems, such as data in DB2 for z/OS.</li></ul>                                                                                                                                                                                                                                                                                                                        |

**Related concepts:**

“IMSplex overview” on page 19

## DB/DC environment

In the DB/DC environment, data is centrally managed for applications that are being executed concurrently and made available to terminal users. Database Recovery Control (DBRC) facilities help to manage database availability, data sharing, and system logging.

The following figure shows an example DB/DC environment.




Figure 2. Example of a DB/DC environment

In the figure, DL/I is shown as part of the control region, but it does not need to run there. You can run DL/I in its own address space.

To understand the DB/DC Environment and how it works, it is important to understand the *transaction*. The basic unit of work in a DB/DC environment is the transaction. Transaction processing consists of:

- Receiving a request for work that has been entered at a terminal. The request is in the form of a transaction code, which identifies the kind of work to be performed and the data needed to do it.
- Invoking a program to do the work, and preparing a response for the terminal operator (for example, an acknowledgment of work performed or an answer to an inquiry).
- Transmitting the response to the terminal that requested the work.

The simplest kind of transaction involves two messages: an input message from the terminal user and an output message in return. Application programs can also send messages to terminals other than the input source, and they can generate transactions.

#### Related concepts:

“Transaction flow in DB/DC and DCCTL environments” on page 384

## Control region

The IMS control region holds the *control program*, which continuously runs in the control region address space and controls the processing in other regions.

The IMS control region services all DL/I calls either directly or through the DL/I separate address space (DLISAS). The IMS control region owns all the databases that can be accessed by online application programs and is responsible for all physical input/output to the databases. The IMS control region supervises the processing of messages and all the communication traffic for the connected terminals. The IMS control region also manages information for restart and recovery purposes and operates the IMS system log.

The IMS control region is normally started by using the z/OS START command. The IMS control region then automatically initiates the DBRC address space.

## Database Recovery Control facility

The Database Recovery Control facility (DBRC) helps you control log and database recovery. It also controls the data sharing environment by allowing or preventing access to databases by the IMS systems that share those databases.

DBRC runs in its own address space, but is subordinate to the IMS control region. DBRC is required by all online IMS systems and any system that uses data sharing.

### Related concepts:

Chapter 29, “Overview of DBRC,” on page 479

## Dependent regions

The dependent regions are separate address spaces from the control region, which are dependent on IMS and in which IMS schedules the applications that process transactions. Dependent regions are initiated by a z/OS START command or by a /START REGION command from the IMS master terminal.

The following are different types of dependent regions:

- Message Processing Program (MPP) Region.  
MPP regions are started either by the master terminal operator or by JCL if the control program is running. The control program schedules application programs within the MPP regions. The application programs then run, accessing the online databases and obtaining their transaction input from the message queues. The application programs cannot access z/OS files or issue z/OS checkpoints. The application program output messages can be directed to LTERMs or to other application programs. An application program can remain scheduled in an MPP region even when there is no work to process for that region. The MPP region remains in a wait-state (wait-for-input mode) until there is more work for the region to process.
- Batch Message Processing (BMP) region.  
z/OS schedules the BMP regions. The application programs in those regions are determined by the JCL used to start each region, not by the control region. These application programs can access databases owned by the control region and z/OS data sets owned by their BMP regions. z/OS data sets include data entry databases (DEDBs) and main storage databases (MSDBs).  
Application programs in BMP regions can access input and output message queues; they can also execute in wait-for-input mode. To access the input message queues, you specify, in the JCL for a BMP region, a transaction code you want to access. Specifying this transaction code also gives you access to the output message queues using terminal program communications blocks (PCBs).

in the application program's specification block (PSB). Even without access to input message queues, if you specify an output LTERM or transaction code in the JCL for the region, the application program can issue output messages.

- IMS Fast Path (IFP) region.

Two types of programs run in IFP regions:

- Application programs for processing Fast Path messages; these are called message-driven programs.
- Utilities that process DEDBs; these are BMPs.

- Java™ Message Processing (JMP) region.

JMP regions process messages with either applications written in Java or applications written in both Java and OO COBOL.

- Java Batch Processing (JBP) region.

JBP regions process batch operations with either applications written in Java or applications written in both Java and OO COBOL.

#### **Related reading:**

- For more information on the types of Fast Path processing and the administration of a DB/DC environment that includes Fast Path, see *IMS Version 11 Database Administration*.
- For more information on writing applications that can run in JMP and JBP regions, see *IMS Version 11 Application Programming*.

### **The master terminal**

The master terminal is the control center of the DB/DC environment. The Master Terminal Operator (MTO) must know all the operating aspects of the system and be familiar with the purpose and action of all the IMS commands that are entered.

Some characteristics of the master terminal are:

- It is used to enter commands that start, stop, and restart the system.
- As a logical terminal, it receives system messages.
- The primary control of the network and connecting terminals is performed through the master terminal. It can start and stop communication lines and assign logical terminals to physical terminal destinations.
- The status of the system can be displayed from the master terminal. Such items as the number of transactions to be processed, the number of programs and databases that are active, and the status of communication lines can be requested.
- If a program or database error occurs, commands can be entered from the master terminal to prevent further processing against the affected resource and to prevent input or output activity for terminals. These recovery actions can also apply to a recovery of the entire system after an abnormal termination occurs.
- If an exception condition occurs, the status of the Online Log Data Set (OLDS) can be displayed, and the function of the OLDS can be controlled from the master terminal.

If the master terminal becomes inoperable, the operating system console can be used as a backup. The MTO can either operate the IMS from the system console or assign the master terminal LTERM (logical terminal) to an alternate terminal. Messages continue to be routed to the old master terminal LTERM until the LTERM is assigned to the system console. The address of the system console is LINE 1 PTERM 1. If the system console is used to continue operation, terminal security is identical to that of the master terminal.



## Databases supported in a DB/DC environment

The DB/DC environment supports all full-function databases (HSAM, SHSAM, HISAM, SHISAM, HDAM, PHDAM, HIDAM, PHIDAM, and PSINDEX).

BMP regions in a DB/DC environment can access GSAM databases. BMP regions can also access external subsystems (for example, DB2 for z/OS), because DB/DC supports the external subsystem interface.

Fast Path data entry databases (DEDBs) and main storage databases (MSDBs) are also supported.

**Note:** PHDAM, PHIDAM, and PSINDEX are database types added for High Availability Large Database (HALDB). They are partitioned equivalents of HDAM, HIDAM, and secondary indexes, respectively.

## Fast DB Recovery region in DB/DC

The Fast DB Recovery region is a separate IMS control region that monitors an IMS, detects failure, and recovers any IMS-owned database resources that are locked by the failed IMS, making them available for other IMS systems without having to wait for the next full restart.

For database resources that are not IMS owned, such as DB2 for z/OS databases, the Fast DB Recovery region provides an optional exit routine, ESAF In-Doubt Notification exit routine (DFSFDN0), for this purpose.

The Fast DB Recovery region is executed by the cataloged procedure that is supplied by IMS system definition. You must start the Fast DB Recovery region after you start the IMS that it tracks.

To enable a DB/DC IMS for Fast DB Recovery, specify the FDRMBR parameter in the IMS procedure. The FDRMBR parameter defines the DB/DC system as Fast DB Recovery-capable.

## Data sharing in a DB/DC environment

Data can be shared among dependent regions and with other IMS systems. The other systems can be DB/DC or DBCTL.

If you intend to share data at the block level, the Internal Resource Lock Manager (IRLM) must be present in every environment that participates. IRLM runs in its own address space.

## Running the Extended Recovery Facility

The Extended Recovery Facility (XRF) is a combination of programs that includes two DB/DC environments to provide a high level of IMS availability to end users.

One environment is active and is called the active system. The other environment continuously tracks the processing of the first and is called the alternate system. The alternate system is ready to take over if the active system fails, or if a planned takeover is initiated (to perform maintenance, for example).

### Related concepts:

Chapter 44, "Extended Recovery Facility Overview," on page 585

## The RSR environment

The remote site recovery (RSR) environment allows you to recover quickly from an interruption of computer services at a primary site.

IMS database and online transaction information is continuously transmitted to a secondary site. In the event of a service interruption at the active site, this secondary site is ready to take over the work from the active site.

**Related concepts:**

Chapter 45, “Remote Site Recovery overview,” on page 663

## DBCTL environment

The DBCTL environment is similar to the DB/DC environment; a DL/I region owns the databases to be processed. DL/I also exists in the DBCTL environment, although DL/I must run in its own address space. Database Recovery Control (DBRC) facilities, required for DBCTL, help to manage database availability, data sharing, system logging, and database recovery.

The greatest dissimilarity between DBCTL and DB/DC is that DBCTL does not support user terminals, a master terminal, or message handling. Therefore, no MPP regions exist. The BMP region is used only by batch applications and utilities. External program subsystems can, however, use an interface that does handle messages—a coordinator controller (CCTL). The same interface exists in the DB/DC control region, so it is possible to use a CCTL with a DB/DC environment. The interface between the CCTL and the control region is the database resource adapter (DRA). The DRA resides in the same address space as the CCTL.

The CCTL handles message traffic and schedules application programs, all outside the DBCTL environment. It passes database calls through the interface to the control region, which sends the calls to DL/I and passes results back through the interface to the CCTL.

The topics in this information that describe the IMS online system applies to both DB/DC and DBCTL. Exceptions are noted as not applicable to DBCTL.

The following figure shows an example of the DBCTL environment.




Figure 3. Example of a DBCTL environment

Even though the DBCTL environment includes no master terminal, you can still control the environment with IMS commands. The commands and command

functions that control message processing are not operable here, but the remainder are. They can be entered through the MVS™ console or a secondary console. The control region recognizes commands by their first character, a slash (/). You can choose a different first character during system definition, or as an execution parameter.

**Note:** References to the master terminal operator (MTO) in this information refer to either the DB/DC MTO or to the DBCTL operator.

Output messages from a command are sent to the console that entered the command. You can also specify other consoles to receive unsolicited output. These consoles are those that fall into the category you define using the IMS-generating macro, IMSCTRL.

### **Databases supported by DBCTL**

The DBCTL environment supports all full-function databases (HSAM, SHSAM, HISAM, SHISAM, HDAM, PHDAM, HIDAM, PHIDAM, and PSINDEX).

BMP regions in a DBCTL environment can access GSAM databases. BMP regions can also access external subsystems (for example, DB2 for z/OS), because DBCTL supports the external subsystem interface.

The DBCTL environment supports Fast Path data entry databases (DEDBs).

**Note:** PHDAM, PHIDAM, and PSINDEX are database types added for High Availability Large Database (HALDB). They are partitioned equivalents of HDAM, HIDAM, and secondary indexes, respectively.

### **Fast DB Recovery region in DBCTL**

The Fast DB Recovery region is a separate IMS control region that monitors an IMS, detects failure, and recovers any IMS-owned database resources that are locked by the failed IMS, making them available for other IMS systems without having to wait for the next full restart.

For database resources that are not IMS owned, such as DB2 for z/OS, the Fast DB Recovery region provides the optional ESAF Indoubt Notification exit routine (DFSFDN0) for this purpose.

The Fast DB Recovery region is executed by the cataloged procedure supplied by IMS system definition. You must start the Fast DB Recovery region after you start the IMS that it tracks.

To enable a DBCTL subsystem for Fast DB Recovery, you specify the FDRMBR parameter in the DBC procedure. The FDRMBR parameter defines the DBCTL system as Fast DB Recovery-capable.

### **Data sharing**

As in the DB/DC environment, data can be shared between dependent regions and with other IMS systems. The other systems can be either DB/DC or DBCTL environments.

If you intend to share data at the block level, IRLM must be present in every environment that participates.

## Alternate DBCTL environment

You cannot have XRF in a DBCTL environment, but you can still run two DBCTL environments—an active and an alternate—and thereby increase system availability.

However, the alternate DBCTL environment does not track the processing of the active environment. The console operator must use the emergency restart command (/ERESTART) against the alternate system in order to make it the active environment.

## DCCTL environment

DCCTL is an IMS Transaction Manager subsystem that has no database components. A DCCTL environment is similar to the DB/DC environment. The primary difference is that a DCCTL control region owns no databases and does not service DL/I database calls.

DCCTL, in conjunction with the IMS External Subsystem Attach Facility (ESAF) or the DB2 for z/OS Recoverable Resource Manager Services Attach Facility (RRMS), provides a Transaction Manager facility to external subsystems (for example, DB2 for z/OS). In a DCCTL environment, transaction processing and terminal management is identical to transaction processing and terminal management in a DB/DC environment. DCCTL contains the programming support necessary for:

- Master terminal support
- Terminal network support
- Data communication
- Message handling
- Transaction processing
- Application program execution
- IMS command execution

DCCTL also supports online change, Message Format Service (MFS), Multiple Systems Coupling (MSC), and Database Recovery Control (DBRC).

All IMS commands are supported in a DCCTL environment except database commands and database-related keywords.

DBRC is required, and is used to maintain system log information for restart. DBRC in a DCCTL environment maintains logs only for transactions. External database subsystems must maintain their own database logs.

DCCTL consists of three address space types:

- Control region
- DBRC
- Dependent regions (up to 999)

Dependent regions and DBRC are subordinate to the control region.

The DCCTL control region contains three structural components:

- A data communication manager, which controls terminal states and input/output message traffic. It also contains security controls that prevent unauthorized access to DC resources.

- A message manager, which is the read/write and I/O interface between terminal input from the data communication manager and the scheduling services of the Transaction Manager.
- A Transaction Manager, which manages MPPs, BMPs, and IFPs, schedules application programs in those dependent regions, and owns and responds to the application programming interface (API).

Each manager (data communication manager, Transaction Manager, and message manager) controls the use of its resources and the recoverability of its resources during a system failure. Like DB/DC dependent regions, MPP, BMP, IFP, JMP, and JBP dependent regions are used by the Transaction Manager to schedule application programs.

The following figure represents a DCCTL environment that is attached to an external subsystem.




Figure 4. Example of a DCCTL environment and attached subsystem

DCCTL coordinates the sync point recovery process with the connected external subsystems. DCCTL ensures that database updates and terminal messages are committed when an application program reaches a sync point.

Two methods exist for connecting a DCCTL control region to another subsystem. You can have DCCTL use the control region EXEC parameter, SSM, to select the PROCLIB member. Or, you can use the /START SUBSYSTEM SSM command, which allows DCCTL to connect to other subsystems even though you did not request this option when you started IMS.

You can also specify the dependent region EXEC parameter, SSM, for dependent regions. The control region SSM member definition allows the dependent region to select one or more external subsystem connections. The SSM member can contain no definitions (null members) to prevent a connection to an external subsystem.

After you use the /START SUBSYSTEM SSM command, you must stop active dependent regions and then restart them if they require an external subsystem connection.

After IMS has established a connection between the dependent region and the external subsystem, a thread is created between the connected regions. The thread is used for subsequent application program calls, committing the data, or in failure situations, backing out the data.

The application programs managed by DCCTL are identical to those managed by the DC manager and TM manager in the DB/DC environment.

#### **Related reading:**

- For more information about the ESAF and RRSF interfaces, see Chapter 28, “Accessing external systems from within IMS,” on page 475.
- For more information on using the /START SUBSYSTEM SSM command, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

### **Databases supported by DCCTL**

The DCCTL configuration of IMS supports and is a compatible communications front end for database and dependent region combinations.

The following database and dependent region combinations are compatible communications front end for a DCCTL configuration of IMS:

- GSAM databases for BMP regions
- DB2 for z/OS databases for BMP, MPP, and IFP regions through the External Subsystem Attach Facility (ESAF)
- DB2 for z/OS databases for JMP and JBP regions through the DB2 Recoverable Resource Services attachment facility (RRSAF)

**Restriction:** DCCTL does not support the following database types:

- Fast Path databases
- Full-function databases

full-function or Fast Path databases.

With GSAM databases, DCCTL uses sequential, non-IMS data sets with a BMP. Application programs can also issue symbolic checkpoint (CHKP) and extended restart (XRST) calls against a GSAM data set using the I/O PCB. The ability to issue CHKP and XRST calls allows data set repositioning.

When DCCTL accesses DB2 for z/OS databases through the ESAF or RRSF interfaces, the control region initiates contact with other subsystems. The other subsystems that DCCTL can access are defined in a subsystem member (SSM) of

the IMS.PROCLIB data set that you provide. The subsystem definition contains the information DCCTL uses to communicate with the other subsystems.

**Related reading:**

- For more information about accessing external subsystems from IMS, see Chapter 28, “Accessing external systems from within IMS,” on page 475.
- For more information about RRSAF, see *DB2 for z/OS Application Programming and SQL Guide*.

DCCTL handles transaction management for online IMS applications that need to access external subsystems.

**Application calls supported by DCCTL**

Application program calls passed to DCCTL receive an AD status code if the call function is not supported or if a database PCB is passed as part of the call list.

The following data communication calls are available to application programs in a DCCTL environment:

- AUTH
- CHNG
- CMD
- GCMD
- GN
- GU
- ICAL
- ISRT
- PURG
- SETO

The following system service calls are available to application programs in a DCCTL environment:

- APSB
- CHKP
- DPSB
- INIT
- INQY
- LOG
- ROLB
- ROLL
- ROLS
- SETS
- SETU
- SYNC
- XRST

**DCCTL compared to DB/DC**

DCCTL and the TM part of a DB/DC environment are very similar.

Similarities include:

- Control region and dependent region initialization and termination.

- System definition and generation.
- Restart.
- Stage 1 input without removing the database definitions. You will need to make some changes to define a DCCTL system. See *IMS Version 11 System Definition* for more information about analyzing macros for system definition.
- Diagnosis.

### **Exit routines**

You do not need to change your IMS exit routines or your existing IMS application programs that access other subsystem resources in a DCCTL environment. However, application programs that contain a mixture of calls that access other subsystems and IMS databases require changes. All DL/I calls that use a database PCB receive status code AD.

### **Automated operator transactions**

Automated operator transactions are started the same way as IMS transactions are started. Automated operator transactions run as IMS application programs with the authority to issue a subset of DCCTL commands using DL/I calls.

## **DB batch environment**

The batch environment consists of a batch region—one address space—where an application program and DL/I routines reside. The batch job that runs here is initiated with JCL, like any operating-system job.

The following figure represents a batch environment in which the batch job is submitted through a TSO terminal, and an application program is run in order to read from an update file, write to a database, and produce a report. For example, an inventory application might be run—one that reads sales records (inventory reductions) and supply records (inventory increases), updates a database accordingly, and prints an inventory sales report.






Figure 5. Example of a DB batch environment

## TM batch environment

IMS TM supports a batch region for running application programs. IMS applications cannot use the ESAF to issue SQL calls in batch. This support is provided by an external subsystem.

You can connect DB2 for z/OS in an IMS TM batch environment in one of two ways. You can use the SSM parameter on the TM batch-region execution JCL and specify the actual name of the batch program on the MBR parameter. Alternatively, you can code the DDITV02 DD statement on the batch-region execution JCL and specify the name of the DB2 for z/OS module, DSNMTV01, on the MBR parameter.

Valid TM batch region types are DBB, DLI, or UPB. All other region types are not applicable to the TM batch environment.

You specify generated program specification blocks (GPSBs) for a TM batch environment using the PSB parameter in the DBBBATCH and DLIBATCH procedures.

The IEFORDER log DD statement is required in order to enable log synchronization with other external subsystems, such as DB2 for z/OS.

### Related reading:

- If your external subsystem is DB2 for z/OS, see the *DB2 for z/OS Application Programming and SQL Guide* for a description of the steps required to allow batch programs to issue SQL calls.
- For additional options or requirements, see *IMS Version 11 System Definition*.

## IMSplex overview

An IMSplex is made up of IMS and z/OS components that work together. Connected systems such as Multiple Systems Coupling and replicated systems such as XRF and RSR are used in an IMSplex.

An IMSplex can be defined as:

- A set of IMS systems working together to share resources or message queues (or both) and workload.
- A single IMS system using the Common Service Layer (CSL) without the Resource Manager (RM) to have a single point of control. This configuration allows you to use IMS type-2 commands, which can be issued only through the Operations Manager (OM) APIs by an automated operator program (AOP). Optionally, IMSplexes of this type can also include the CSL Open Database Manager (ODBM).
- A single IMS system using the CSL with an RM.

Contrast the description of an IMSplex with a description of a Parallel Sysplex<sup>®</sup> (hereafter referred to as sysplex). A sysplex is multiple z/OS images working together, connected by a coupling facility. One or more IMSplex systems can be defined on one or more z/OS images; however, you do not have to have an IMS instance on every z/OS image in the sysplex.

The concept of data sharing is important in an IMSplex, specifically in an IMSplex that shares resources and workload.

This allows IMS batch jobs, online applications, or both, to run anywhere in the IMSplex. These batch jobs and applications can also access all of the data in the shared IMS databases. You can distribute your IMS workload according to your business requirements. After data sharing is enabled, you must create a plan for distributing batch and transaction workload across the IMSplex and then implement that plan. One technique is using shared queues (SQ). With SQ, any transaction can be entered from the network to any IMS system in the IMSplex and be executed by any IMS system in a shared-queues group. If one IMS system is unavailable, a different system can handle the work.

If you have multiple IMS systems sharing resources or message queues, administration and operation can be simplified by using the IMS Common Service Layer (CSL), which reduces complexity by providing a single-image perspective. With CSL, you can manage multiple IMS systems as if they were only one system. For example, instead of entering type-1 commands from multiple MTOs or a single point of control (SPOC) to perform local online changes, you can enter a type-2 command from one SPOC to perform global online change. The following figures contrast the relative complexity of managing an IMSplex with and without a CSL.

Managing an IMSplex using CSL provides:

- Improved systems management
- A single-system image
- Simpler operations through a single point of control
- Shared resources across IMS systems




Figure 6. Complex management of an IMSplex without CSL




Figure 7. Simplified management of an IMSplex with CSL

When you manage an IMSplex with the CSL, the resulting environment is comprised of multiple *IMSplex components*. An *IMSplex component* either manages resources, manages operations, or facilitates communication among *IMSplex components*. The following are *IMSplex components*:

- Operations Manager (OM)
- Resource Manager (RM)
- Open Database Manager (ODBM)

- Structured Call Interface (SCI)
- IMS Connect
- Database Recovery Control (DBRC)
- Common Queue Server (CQS)
- TSO Single Point of Control (SPOC)
- REXX SPOC API for automation programs
- Online control regions

Any program that registers with the SCI is considered an IMSplex component. Users and vendors can also write programs that register with SCI; these programs are also considered IMSplex components. Note that batch regions and dependent regions (MPR, BMP, and IFP) are not IMSplex components. Although DBRC code running within a batch region registers with SCI, the batch region itself is not a component. When an IMSplex component is initialized and joins the IMSplex, it becomes an *IMSplex member*.

#### **Related concepts:**

“IMS environments and configurations” on page 5

### **z/OS components and system services used in an IMSplex**

The z/OS components that are used in an IMSplex include the Coupling Facility and its cache, list, and lock structures.

An IMSplex takes advantage of several z/OS system services:

#### **z/OS cross-system coupling facility**

An IMSplex uses XCF for communication services. XCF services allow authorized programs on one system to communicate with programs on the same system or on other systems.

#### **automatic restart management (ARM)**

An IMSplex uses ARM for recovery services. ARM restarts subsystems, all the elements of a workload on another z/OS image after a z/OS failure, and a failed z/OS image.

#### **cross-system extended services (XES)**

An IMSplex uses XES for data sharing services and system-managed rebuild support.

### **IMS components and system services that are part of an IMSplex**

An IMSplex can include several different IMS components: IMS subsystems, DBRC and RECON data sets, IMS Connect, and TSO SPOC. IMS uses IRLM for lock services.

In addition, a number of IMS address spaces provide system services in an IMSplex. These address spaces are all built using the Base Primitive Environment (BPE), which provides system services such as tracing, message formatting, parsing, storage management, sub-dispatching, and serialization.

The BPE system services are used by the Common Queue Server (CQS), and the CSL's Open Database Manager (ODBM), Operations Manager (OM), Resource Manager (RM), Structured Call Interface (SCI), and, optionally, the Database Recovery Control facility (DBRC), each of which contributes other system services:

#### **Common Queue Server**

Receives, maintains, and distributes data objects from shared queues.

## Common Service Layer

Includes the following components:

### Open Database Manager

Provides access to IMS databases managed by IMS DB systems in DBCTL and DB/TM environments within an IMSplex. ODBM supports TCP/IP clients through IMS Connect and application servers running application programs that use the IMS ODBA interface, such as DB2 for z/OS or WebSphere® Application Server for z/OS.

### Operations Manager

Routes commands, consolidates command responses, provides an API for command automation, and provides user exits for customization and security.

### Resource Manager

Maintains global resource information, ensures resource consistency, and coordinates IMSplex-wide processes such as global online change.

### Structured Call Interface

Routes messages, registers and deregisters members of an IMSplex, provides security authentication of IMSplex members, and more.

IMS also uses IRLM for lock services.

## Multiple Systems Coupling in an IMSplex

The Multiple Systems Coupling (MSC) function allows you to distribute and balance your workload across multiple systems.

MSC distributes transaction and message traffic between IMS subsystems. MSC extends its communication with IMS and scheduling capabilities, enabling you to perceive one virtual IMS system, when in fact transactions are being routed among several IMS subsystems.

An MSC network can coexist with an IMSplex using shared queues. A temporary MSC and IMSplex coexistence can be useful if you are migrating from an MSC network to an IMSplex configuration. Coexistence can also be permanent, such as when an MSC link connects an IMSplex to an IMS system outside of the IMSplex.

In both coexistence examples, you must consider the proper routing and processing of your transaction messages across both the MSC and IMSplex environments, because each environment uses a different routing method. Generally, MSC networks route transactions to specific IMS systems using SYSIDs. IMSplex systems with shared queues route transactions by making them available on the shared queue to any IMS system that registers an interest in the transactions. When an MSC network and an IMSplex with shared queues coexist, both of these methods of routing can apply to transactions.

For more information on the coexistence of MSC systems and IMSplex systems, see *IMS Version 11 Communications and Connections*.

## Recovering systems in an IMSplex

The recovery of an IMS system can be managed by Extended Recovery Facility (XRF) or Remote Site Recovery (RSR). An IMS Fast Database Recovery region (FDBR) can also be used for recovery purposes.

**XRF** Performs takeover on a system-by-system basis; XRF relies on the same physical databases and many of the same physical resources as the active IMS system.

**RSR** Remote takeover includes all IMS systems that share databases. RSR duplicates all physical resources without a distance limitation. No single point of failure exists, so an area-wide physical problem can be overcome.

XRF can be used in an IMSplex in the same way that it is used on a local system. The XRF alternate system needs to have a unique IMSID in the IMSplex, be defined in the IMSplex and have access to the CSL. An SCI must reside on the z/OS image where the XRF alternate resides. You cannot use XRF to recover an entire IMSplex, but you can use XRF to recover an individual IMS in an IMSplex.

With RSR, IMS database and online transaction information is continuously transmitted to a tracking (remote, or secondary) site. The remote site is continually ready to take over the work of the active site in the event of service interruption at the active site. If the active site is an IMSplex with CSL, then the remote site must be a different IMSplex with CSL.

An IMS Database Control (DBCTL) warm standby region and an IMS Fast Database Recovery region (FDBR) can also be used for recovery purposes.

**FDBR** An FDBR region monitors an IMS subsystem and can automatically recover database resources (shared databases and areas) if that IMS subsystem fails.

**DBCTL standby**

While one DBCTL subsystem is active, you can start another DBCTL subsystem as a standby alternate. This alternate subsystem does not track the active subsystem (as it would in an XRF complex), but is a fully initialized IMS DBCTL subsystem that is waiting for a restart command.

An IMSplex can include XRF, RSR, FDBR, and DBCTL standby regions.

**Related reading:**

- For more information on XRF and RSR, see Part 4, “IMS system recovery,” on page 583.
- For more information on FDBR, see *IMS Version 11 Operations and Automation*.

**Defining and tailoring an IMSplex**

You must establish the procedures required to start an IMSplex, and specify the appropriate parameters in the DFSCGxxx member or the CSL section of the DFSDFxxx member of the IMS PROCLIB data set.

The following topics describe how to initialize each of the necessary IMSplex components for an IMSplex with CSL.

These topics briefly describes the components that you must define in an IMSplex.

**Related reading:**

- For detailed information about IMSplex system definition, see *IMS Version 11 System Definition*.
- For more information about defining the IMS control region, the system definition macros, and the PROCLIB member data sets, see *IMS Version 11 System Definition*.

**Related concepts:**

“Global TM resource management” on page 51

**Related tasks:**

“Defining a simplified IMSplex for the type-2 command environment” on page 25

**Defining CQS:**

To define CQS, create the CQS and BPE PROCLIB member data sets and define the execution data sets.

The following steps summarize defining CQS.

1. Create a coupling facility resource management (CFRM) policy for the MSGQ, EMHQ and RSRC structures.
2. Define the applicable z/OS policies.
3. Activate the CFRM policy.
4. Create the CQS and BPE PROCLIB member data sets.
5. Define all CQS execution data sets.
6. Customize your CQS environment.
7. Authorize connections to CQS structures.
8. Update the z/OS program properties table.

**Defining the IMS control region:**

To initialize IMS to function in an IMSplex, you must define the IMS control region to communicate with OM for command processing and RM for resource validation, status updates, and global online change.

Steps in the definition process include:

- Uniquely identifying each IMS system in the IMSplex by using the IMSID= keyword in the system definition macro
- Identifying the IMSplex to each IMS system by using the IMSPLEX= keyword in the DFSCGxxx PROCLIB member data set, or the CSL section of the DFSDFxxx PROCLIB member data set, of each IMS system

The PROCLIB member data sets that must be modified to correctly define the IMS system and the IMSplex include:

- DFSCGxxx
- DFSDFxxx
- DFSPBxxx
- DFSSQxxx
- DFSVSMxx

**Defining ODBM:**

Use the CSLDIxxx member of the IMS.PROCLIB data set to specify parameters for initializing ODBM. Certain CSLDIxxx parameters can be overridden with ODBM execution parameters.

The IMSplex name that you specify in CSLDIxxx by using the IMSPLEX(NAME=*plxnm*) parameter must be the same as the IMSplex names specified in the CSLOIxxx, CSLRIxxx, CSLSIxxx, DFSDFxxx, and DFSCGxxx PROCLIB member data sets.

Use the CSLDCxxx PROCLIB member to specify the parameters that configure the ODBM connections to IMS data stores.

### Defining OM:

Use the CSLOIxxx member of the IMS.PROCLIB data set to specify parameters for defining OM. Certain CSLOIxxx parameters can be overridden with OM execution parameters.

The IMSplex name that you specify in CSLOIxxx by using the IMSPLEX(NAME=*plxnm*) parameter must be the same as the IMSplex names specified in the CSLDIxxx, CSLRIxxx, CSLSIxxx, DFSDFxxx, and DFSCGxxx PROCLIB members.

### Defining RM:

Use the CSLRIxxx member of the IMS.PROCLIB data set to specify parameters for defining RM. Certain CSLRIxxx parameters can be overridden with RM execution parameters.

The IMSplex name that you specify in CSLRIxxx by using the IMSPLEX(NAME=*plxnm*) parameter must be the same as the IMSplex names specified in the CSLDIxxx, CSLOIxxx, CSLSIxxx, DFSDFxxx, and DFSCGxxx PROCLIB members.

### Defining SCI:

Use the CSLSIxxx member of the IMS.PROCLIB data set to specify parameters for defining the SCI address space. Certain CSLSIxxx parameters can be overridden with SCI execution parameters.

The IMSplex name that you specify in CSLSIxxx by using the IMSPLEX(NAME=*plxnm*) parameter must be the same as the IMSplex names specified in the CSLDIxxx, CSLOIxxx, CSLRIxxx, DFSDFxxx, and DFSCGxxx PROCLIB members.

### Defining a simplified IMSplex for the type-2 command environment:

To define a simplified IMSplex without RM either specify RMENV=N in the DFSCGxxx member of the IMS PROCLIB data set or specify RMENV=N in the CSL section of the DFSDFxxx PROCLIB member while defining an IMSplex.

To bring up a CSL without RM, RMENV=N must be specified in the SCIPROC parameter in the DFSCGxxx or DFSDFxxx members of the IMS PROCLIB data set.

#### Related concepts:

“Global online change” on page 49

“A simplified CSL configuration” on page 41

#### Related tasks:

“Defining and tailoring an IMSplex” on page 23

➡ Starting the CSL OM (Operations and Automation)

➡ Starting the CSL SCI (Operations and Automation)

#### Related reference:

“CSL configuration examples” on page 54



➡ DFSCGxxx member of the IMS PROCLIB data set (System Definition)

➡ DFSDFxxx member of the IMS PROCLIB data set (System Definition)

### **Defining the IMSplex name for RECON loss notification or DBRC parallel RECON access:**

To enable automatic RECON loss notification or parallel RECON access, you must set an IMSplex name for the RECON data set using the CHANGE.RECON IMSPLEX command.

If automatic RECON loss notification is the only use of the CSL in an IMSplex, SCI is the only CSL address space that is required.

All DBRC instances accessing this RECON must specify the IMSplex name for the RECON data set by using the DBRC SCI Registration exit routine (DSPSCIX0) or the EXEC statement keyword parameter IMSPLEX=. To use SCI, DBRC must provide the name of the IMSplex in which the DBRC instance is running.

The DBRC SCI Registration exit routine, DSPSCIX0, is an exit routine that DBRC calls before it registers with SCI. If you have enabled parallel RECON access, this exit routine must supply the IMSplex name needed for SCI registration. If the exit routine is not used, DBRC continues as if the sample version of the exit routine was being used, which means that DBRC does not register with SCI. DSPSCIX0 must be found in an authorized library or in LINKLST. In a TSO environment, the library that contains DSPSCIX0 must be a TSO TASKLIB library.

If more than one set of RECONs is defined in the IMSplex, a *group\_id* also needs to be defined either in the DSPSCIX0 exit routine or in the RECON data set. The EXEC statement keyword parameter for the *group\_id* is DBRCGRP=xxxx. To define a *group\_id* in the RECON data set, use the CHANGE.RECON IMSPLEX(*imsplex\_name*,*group\_id*) command.

### **Defining global online change:**

To enable global online change in an IMSplex that is defined with a CSL, the DFSCGxxx member or the CSL section of the DFSDFxxx member of the IMS PROCLIB data set must be specified.

The DFSCGxxx PROCLIB member and the CSL section of the DFSDFxxx PROCLIB member include parameters related to SCI, OM, and RM, which are common to all IMS systems in the IMSplex (except those systems defined without an RM). The OLC= parameter must specify GLOBAL for global online change or LOCAL for local online change.

OLC=GLOBAL means that the online change is coordinated across the IMSplex. Global online change is prepared and committed by the INITIATE OLC command. If global online change is enabled and a resource structure is defined, the MODBLKS, FORMAT, and ACBLIB data sets must be consistent across the IMSplex, unless resource consistency checking is omitted by the NORSCCC keyword.

OLC=LOCAL means the online change applies locally to each IMS. Local online change is prepared and committed by the /MODIFY PREPARE and /MODIFY COMMIT commands on each local IMS. The local online change must be manually coordinated across an IMSplex.

Specify the OLCSTAT data set using the OLCSTAT= parameter of the DFSCGxxx PROCLIB member data set or the CSL section of the DFSDFxxx PROCLIB member data set if global online change is enabled in an IMSplex with CSL. Run the Global Online Change utility (DFSUOLC0) to:

- Initialize the OLCSTAT data set before IMS cold starts for the first time.
- Recreate the OLCSTAT data set if an error makes the OLCSTAT data set unusable.
- Unlock the OLCSTAT data set.

**Restriction:** Global online change is not supported for RSR Tracker. If global online change is specified, the RSR tracking IMS abends.

*OLCSTAT data set:*

The global OLCSTAT data set, which z/OS assigns to IMS, is a basic sequential access method (BSAM) data set that supports one record of variable size. You can allocate and deallocate the OLCSTAT data set dynamically. Unlike the MODSTAT data set, the OLCSTAT data set does not need to be continuously available when your IMS system is running.

### OLCSTAT data set attributes

Before the OLCSTAT can be allocated, you must define the data set attributes. The following table shows the recommended attributes for the OLCSTAT data set, to allow up to 65 IMS records.

*Table 2. Recommended attributes for the OLCSTAT data set*

| Data set attribute | Value      |
|--------------------|------------|
| DSORG              | Sequential |
| RECFM              | V          |
| LRECL              | 5204       |
| BLKSIZE            | 5208       |

If you do not define these attributes, when IMS attempts to allocate the OLCSTAT data set, that allocation fails.

After you define these attributes, you must name the data set by using the OLCSTAT parameter in the DFSCGxxx PROCLIB data set member.

### OLCSTAT data set format

The OLCSTAT data set consists of a header and any number of IMS records, with the following formats:

- OLCSTAT format for VERS=1  

```
hlenhverivermstr  OLCINP modifyid MODBLKSz x IMSACBb x FORMATc x
ilenrveriverims   prepare_timestamp commit_timestamp
```
- OLCSTAT format for VERS=2  

```
hlenhverivermstr  OLCINP modifyid MODBLKSz x IMSACBb x FORMATc x
MOLCstate MOLCtoken MOLCid
ilenrveriverims   prepare_timestamp commit_timestamp
```

The header information includes the master of the last online change phase, an online change in progress lock, the modify ID, the active MODBLKS DD name, the

active IMSACB DD name, the active FORMAT DD name, and the libraries that were changed with the last online change. The header is initialized by the Global Online Change utility (DFSUOLC0).

**hlen** The OLCSTAT data set header record length. The header record is the first line.

**hver** The OLCSTAT data set header version.

**iver** IMS version.

**mstr** IMS ID of the IMS system that was the master of the last online change phase (prepare, commit, or abort).

**ilen** The OLCSTAT data set IMS record length. The OLCSTAT data set contains zero, one or more OLCSTAT data set IMS records.

**rver** The OLCSTAT data set IMS record version.

**iver** IMS version.

**ims** IMS ID.

#### OLCSTAT data set header

The following example shows a sample OLCSTAT data set header for Version 2.

```
RECORD SEQUENCE NUMBER - 1
000000 00000080 00000002 F1F0F1F0 40404040 40404040 40404040 40404040
000020 F0F0F0F1 40D4D6C4 C2D3D2E2 C140D540 C9D4E2C1 C3C2C140 D540C6D6
000040 C140D540 D4D6D3C3 C3D4D7E3 40BF88D7 940BE91E 82402006 283F2321
000060 028D4040 40404040 40404040 40404040 40404040 40404040 40404E40
000080 00000050 00000001 F1F0F1F0 40404040 C9D4E2F1 40F2F0F0 F6F2F8F3
0000A0 F0F5F2F7 4060F0F7 7AF0F040 F2F0F0F6 F2F8F340 F2F3F2F1 F0F3F440
0000C0 F0F04040 40404040 40404040 4040404E
```

```
F0F0F0F0 *.....1010 0000*
D9D4C1E3 *0001 MODBLKSA N IMSACBA N FORMAT*
03465440 *A N MOLCCMPT ..P..Z.. ..... *
40404040 *.. + *
40F2F3F2 *...&....1010 IMS1 2006283 232*
60F0F77A *0527 -07:00 2006283 2321034 -07:*
          *00 + *
```

```
IDC0005I NUMBER OF RECORDS PROCESSED WAS 1
IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0
IDCAMS SYSTEM SERVICES
```

TIME: 1

Each IMS record contains information about an IMS that is current with the online change libraries. The IMS either participated in the last online change, or cold started since the last online change. If no IMS has cold started yet, the OLCSTAT data set contains no IMS records.

You must also run the Online Change Copy utility (DFSUOCU0) as one step in the process of preparing an IMS or an IMSplex for an online change. The DFSUOCU0 utility copies a source library with your new definitions to a target library. In an IMSplex where IMS systems are not cloned and the libraries are not shared, you might need to run the DFSUOCU0 utility on every IMS system in the IMSplex. In an IMSplex where IMS systems are cloned and the libraries are shared, you might need to run the DFSUOCU0 utility only once on one IMS.SDFSRESL data set at the highest IMS level.

The DFSUOCU0 utility can copy the contents of the staging library to a target library that is specified as a parameter. The utility supports a target library value of G, which means that the library target is the active library determined by the utility, using the OLCSTAT data set. The target is the library that is not currently in use by the IMS systems in the IMSplex. If the new target library G is specified, the DFSUOCU0 utility reads the OLCSTAT data set to determine the target library. The global online change commands, INITIATE OLC PHASE(PREPARE), followed by INITIATE OLC PHASE(COMMIT), change the inactive library to an active library.

**Related reference:**

➡ Global Online Change utility (DFSUOLC0) (System Utilities)

➡ Online Change Copy utility (DFSUOCU0) (System Utilities)

### **Dynamic resource definition in an IMSplex**

An IMSplex can include IMS systems that have DRD enabled and others that do not have DRD enabled. This kind of IMSplex environment is referred to as a mixed-DRD IMSplex.

If a command that contains the MODBLKS parameter is issued in a mixed-DRD IMSplex, the MODBLKS parameter is ignored.

Two scenarios are provided to demonstrate how to add application programs in a mixed-DRD IMSplex.

#### **Scenario 1: add an application program to a mixed-DRD IMSplex**

In this IMSplex, IMS1 has DRD enabled, and IMS2 and IMS3 do not have DRD enabled.

1. Perform an ACBGEN for the PSB into the staging ACBLIB
2. Perform a MODBLKS generation to define the APPLCTN program in the staging MODBLKS data set for the IMS2 and IMS3 systems.
3. Issue the CREATE PGM command on IMS1 to dynamically create the program.
4. Copy the staging MODBLKS data set and ACBLIB to the inactive MODBLKS data set and ACBLIB.
5. Perform global online change for the MODBLKS data set and ACBLIB by issuing the following commands:  

```
INITIATE OLC PHASE(PREPARE) TYPE(ACBLIB,MODBLKS)
INITIATE OLC PHASE(COMMIT)
```

The MODBLKS keyword is ignored by the IMS1 system; however, the IMSplex synchronizes and adds the program with one online change instance. The command that includes the MODBLKS keyword is not rejected; this would require another online change for the ACBLIB on the IMS1 system.

#### **Scenario 2: add an application program and its MFS format in a mixed-DRD IMSplex that includes DBCTL IMS systems**

IMS1 and DBCTL1 have DRD enabled. IMS2 and DBCTL2 do not have DRD enabled.

1. Perform an ACBGEN for the PSB into the staging ACBLIB.
2. Perform a MODBLKS generation to define the APPLCTN program in the staging MODBLKS data set for the IMS2 and DBCTL2 systems.
3. Run the MFS utility to create an MFS format in the staging FMTLIB for the IMS1 and IMS2 systems.

4. Issue the CREATE PGM command on the IMS1 and DBCTL1 systems to dynamically create the application program.
5. Copy the staging MODBLKS data set and ACBLIB to the inactive MODBLKS data set and ACBLIB.
6. Perform global online change for MODBLKS, FMTLIB, ACBLIB by issuing the following commands:  

```
INITIATE OLC PHASE(PREPARE) TYPE(ACBLIB,FMTLIB,MODBLKS)
INITIATE OLC PHASE(COMMIT)
```

The MODBLKS keyword is ignored by the IMS1 and DBCTL1 systems. The FMTLIB keyword is ignored by the DBCTL1 and DBCTL2 systems. The ACBLIB keyword is processed on all four systems. This synchronizes the IMSplex and successfully adds the application program and its MFS format with one online change instance. The command that includes the MODBLKS keyword is not rejected; this would require additional online changes on IMS2, DBCTL1, and DBCTL2.

**Related reading:** For a detailed description of what you can do with DRD, see *IMS Version 11 System Definition*.

### **Establishing IMSplex security**

The way you protect the IMS online system does not change for an IMS system in an IMSplex. The resources that can be protected and the facilities available to protect them are similar. The types of security issues that exist for a standalone IMS system also exist for an IMS system within an IMSplex.

The basic decisions you must make when implementing security for IMS include:

- Determine what resources need to be protected.
- Determine which users need access to the resources.
- Determine the level of access to a resource that the users need.

In addition to the security checks made for standalone IMS systems, an IMSplex with a CSL performs additional security checks.

---

## Chapter 2. Concepts for the system administrator

This topic presents concepts that are central to controlling the resources of an IMS. The other topics included in this information assume a thorough understanding of these concepts.

---

### System support for application programs

Before you can run an application program under IMS, control blocks must be defined, generated, and placed into the following libraries: IMS.DBDLIB, IMS.PSBLIB, and IMS.ACBLIB.

Before executing an application program under IMS, you must:

- Describe that program and its use of logical terminals and logical data structures through a program specification block (PSB) generation, using the PSB Generation utility. The PSBs are maintained in one or more IMS system libraries called a PSB library.
- Create a database descriptor block (DBD) to have access to an IMS database. The DBD is assembled into a system library called a DBD library.
- Combine and expand the PSB and DBD control blocks into an internal format called application control blocks (ACBs). The Application Control Blocks Maintenance utility is used to create the ACBs and they are placed in an ACB library.

Application programmers and database administrators need to know the naming conventions for the PSB, DBD, and ACB library data sets.

The system administrator must decide:

- Whether to have the ACB library allocated by JCL or dynamically (using the DFSMDA macro)  
To begin using the DFSMDA macro for dynamic allocation of ACB library data sets, complete the following steps:
  1. Create DFSMDA members for the ACBLIBA and ACBLIBB data sets.  
DFSMDA members can be placed in either the data set specified in the IMS STEPLIB concatenation or in the IMSDALIB DD statement.
  2. Remove the IMSACBA and IMSACBB DD statements from the IMS and DL/I JCL procedures.
  3. Stop IMS and then restart it with the DFSMDA members.
- Whether to have the ACBs loaded into 64-bit storage  
At execution time, the ACBs needed by an application program must be in the 31-bit non-resident ACB storage pool. At scheduling time, IMS first searches the non-resident storage pools to see if the ACBs are already there. If they are not in the pool, IMS can load them into the 31-bit non-resident ACB storage pool from either the ACB library or from a 64-bit storage pool.  
To enable the use of a 64-bit storage pool for ACBs, complete the following steps:
  1. Calculate how much storage to allocate for the 64-bit storage pool by multiplying the total number of non-resident ACB members in the ACB library times the size of these members.

2. Specify the size of the 64-bit ACB storage pool (in gigabytes) on the ACBIN64 parameter in the DATABASE section of the DFSDFxxx PROCLIB member. For most installations, specifying one or two gigabytes for the 64-bit ACB storage pool should be sufficient.

3. Stop IMS and then restart it with the ACBIN64 parameter.

After enabling the use of a 64-bit ACB storage pool, IMS populates this pool the first time an application program that need the ACBs is scheduled. Thereafter, when an application program is scheduled that needs these ACBs, IMS reads them from the 64-bit storage pool instead of from the ACB library.

If you need to remove ACBs from both the 31-bit non-resident storage pool and the 64-bit storage pool, use the DELETE DB and DELETE PGM commands. You cannot remove ACBs from the pools using any command.

**Related reading:** For an expanded overview of generating PSBs, DBDs, and ACBs, see *IMS Version 11 Database Administration*.

**Related reference:**

“XRF requirements” on page 595

---

## Dynamic allocation with IMS

If data sets are specified with JCL in a control region procedure, they are initially allocated when the control region starts up. Some data sets can also be dynamically allocated.

You can specify that the following data sets be dynamically allocated when needed and deallocated when no longer in use.

Using the IMS macro DFSMDA, you declare those data sets that are subject to dynamic allocation and deallocation.

- Database data sets can be dynamically allocated explicitly with the /START command or implicitly when an application program is scheduled. Database data sets can be deallocated with the /DBRECOVERY command.
- For DEDB area data sets, an implicit allocation occurs at first access by an application program; the /STOP command also deallocates the data set.
- An IMS Monitor data set can be dynamically allocated at the time the IMS Monitor is started with the /TRACE SET ON command and deallocated by the /TRACE SET OFF command.
- RECON data sets, online log data sets (OLDSs), write-ahead data sets (WADSs), and system log data sets (SLDSs) that are required as input to restart can be dynamically allocated.
- ACB libraries can be dynamically allocated.
- For High Availability Large Databases (HALDBs), dynamic allocation allocates only the DBRC registered data sets. Allocation does not look for or process any DFSMDA members. The DD names allocated for HALDBs contain the letters A through J, X, or L, suffixed to the 7-byte HALDB partition name.

If an allocation already exists with a DD name that matches the HALDB partition DD name generated, the data set name is compared to the DBRC registered data set name in online IMS environments. If the data set names do not match, an allocation failure occurs.



All data sets using dynamic allocation must be cataloged, except an IMS Monitor data set, which must not be cataloged. A data set that is initially allocated with JCL can be dynamically deallocated and reallocated during the execution of the control region.

**Related reading:** For more information on the IMS macro DFSMDA, see *IMS Version 11 System Definition*.

---

## Base Primitive Environment overview

The IMS Base Primitive Environment (BPE) is a common system service base upon which many other IMS components are built. BPE provides services such as tracing, message formatting, parsing, storage management, sub-dispatching, and serialization.

The following IMS components use BPE:

- Common Queue Server (CQS)
- IMS Connect
- Open Database Manager (ODBM)
- Operations Manager (OM)
- Resource Manager (RM)
- Structured Call Interface (SCI)

You can optionally have Database Recovery Control facility (DBRC) use BPE.

When an IMS component that uses BPE is started, the component loads a copy of the BPE service modules into its address space from the IMS program libraries. The IMS component's modules are specific to that component; however, the BPE service modules are common across the various address spaces. The base system service functions are therefore identical in every address space that uses BPE. The following figure shows the relationship of BPE, IMS components, and IMS program libraries.






Figure 8. BPE and IMS components

Most of the time, BPE is a hidden layer in an IMS component address space. However, you can use the following external interfaces to BPE:

### Configuration

You can configure certain attributes about an address space that uses BPE at startup by using statements in BPE PROCLIB member data sets. For example, you can set the default level and size for BPE-managed trace tables. Information about defining BPE with the BPE PROCLIB member data set is in *IMS Version 11 System Definition*.

### Commands

You can use the small set of commands that BPE provides to operate on BPE-managed resources. For example, you can display and change attributes of BPE-managed user exit routines and trace tables. Information about BPE commands is in *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*.

### User exit routines

You can customize the operation of an IMS component address space by writing user exit routines. Components running with BPE can use the BPE user exit routine service to call component-specific user-supplied exit routines. BPE also has user exit routines of its own. User exit routines that are called through BPE receive control with a standard BPE user exit routine interface and are allowed to use BPE user exit routine callable services. Information about BPE user exits is in *IMS Version 11 Exit Routines*.

## Messages and Abends

BPE has its own messages and abend codes. *IMS Version 11 Messages and Codes, Volume 4: IMS Component Codes* includes information about BPE abend and service return codes and BPE messages.

## Tracing BPE components

You can enable tracing for BPE components (BPE, CQS, DBRC, IMS Connect, ODBM, OM, RM, or SCI), and control whether the trace information is written either to memory, or to both memory and an external trace data set.

This topic describes the tasks associated with enabling tracing for BPE components:

### Enabling BPE tracing

To enable BPE tracing so that it is always on for an IMS address space, you must define or modify the TRCLEV statements in the BPE configuration parameter member of the IMS.PROCLIB data set (specified by BPECFG=). You can also dynamically start BPE tracing by using the UPDATE TRACETABLE command.

The BPE configuration parameter member defines the BPE execution environment settings for the address space being started. BPE trace records can be written to internal (memory only) trace tables and to external data sets.

The default is to write to internal trace tables (EXTERNAL=NO). To write to an external data set, you must set EXTERNAL=YES on the TRCLEV statement and specify the EXTTRACE parameter in the BPE configuration parameter member. You must also define a generation data set group (GDG) for BPE to trace into. If you specify EXTERNAL=YES, trace data is written to both an external data set and internal trace tables.

You can dynamically change the external trace data set specification in the BPE configuration PROCLIB member and refresh the member while an address space is running. For example, if you are running without an external trace data set, you can edit your BPE PROCLIB member, add an external trace data set specification, and start using external trace without having to restart the address space.

#### Related reference:

 BPE UPDATE TRACETABLE command (Commands)

#### Writing to internal trace tables:

The default BPE execution environment settings for the address space is to write to the internal trace table (EXTERNAL=NO).

To write BPE trace records to internal trace tables, specify the type, level, and, optionally, the number of storage pages allocated for the trace table on the TRCLEV parameter of the BPE configuration parameter member (BPECFG=) of the IMS.PROCLIB data set.

#### Related reference:

 BPE configuration parameter member of the IMS PROCLIB data set (System Definition)

### Writing to external data sets:

To write to an external data set, you must set EXTERNAL=YES on the TRCLEV statement and specify the EXTTRACE parameter in the BPE configuration parameter member.

To write BPE trace records to an external data set, you must:

1. Define the external trace data set Generation Data Group (GDG).
2. Define the GDG model.
3. Specify the EXTTRACE parameter in the BPE configuration parameter member (BPECFG=) of the IMS.PROCLIB data set. The EXTTRACE statement should specify the name of the GDG you defined in step 1.
4. Specify which trace tables you want to write to the BPE external trace data sets. You can do this by using either or both of the following methods:
  - Include the EXTERNAL=YES parameter with the other TRCLEV options on the TRCLEV statements for the trace tables you want to externalize. The TRCLEV statements are coded in the BPE configuration parameter member (BPECFG=) of the IMS PROCLIB data set.

**Tip:** Changes to TRCLEV statements are only processed when an address space is started. If your address space is running when you make the changes, they will not take effect until the address space is restarted.

- Issue the UPDATE TRACETABLE command, specifying the trace tables you want to externalize on the NAME parameter, and specifying EXTERNAL(YES). This command is used while an address space is running to dynamically turn external tracing on, regardless of the TRCLEV parameter specification.
5. Optional: If you want to start tracing to external data sets but you did not specify the EXTTRACE parameter when you started the address space:
    - a. Add the EXTTRACE data set to the BPE configuration member of the IMS.PROCLIB data set while the address space is running.
    - b. Issue the UPDATE TRACETABLE command and specify the OPTION(REREAD) option to cause the BPE to process the newly-updated EXTTRACE parameter.

Information about the EXTTRACE and TRCLEV statements and BPECFG= and the IMS.PROCLIB data set is in *IMS Version 11 System Definition*.

### Defining an external trace data set:

Use the z/OS DEFINE GENERATIONDATAGROUP command to define the GDG. The value that you specify for the NAME keyword in this command must match the data set name in the GDGDEF parameter of the EXTTRACE statement of the BPE configuration parameter member (BPECFG=) of the IMS PROCLIB data set.

### Defining a GDG from an IDCAMS job

You can issue the z/OS DEFINE GENERATIONDATAGROUP command from an IDCAMS job, as shown in the following example.

```
//DEFGDG JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
        DEFINE GENERATIONDATAGROUP -
```

```
(NAME(BPEEXTRC.GDG01) -
NOEMPTY -
SCRATCH -
LIMIT(255))
```

### Defining a GDG from a TSO session

You can issue the z/OS DEFINE GENERATIONDATAGROUP command from a TSO session, as shown in the following example.

```
DEFINE GENERATIONDATAGROUP (NAME(BPEEXTRC.GDG01) NOEMPTY SCRATCH LIMIT(255))
```

### Defining a GDG from a TSO batch job

You can issue the z/OS DEFINE GENERATIONDATAGROUP command from a TSO command batch job, as shown in the following example.

```
//DEFGDG      JOB    ...
//STEP1       EXEC  PGM=IKJEFT01
//SYSTSPRINT DD   SYSOUT=A
//SYSTSIN     DD    *
DEFINE GENERATIONDATAGROUP (NAME(BPEEXTRC.GDG01) NOEMPTY SCRATCH LIMIT(255))
```

For more information about the DEFINE GENERATIONDATAGROUP command, see *z/OS DFSMS Access Method Services for Catalogs*.

For additional information on trace data, see *IMS Version 11 Diagnosis*.

#### Defining the GDG model data set:

You must provide a data control block (DCB) to be used when GDG entries are created. You can provide it when you define the generation data group (GDG) model data set. To define the GDG model data set, use the z/OS program IEFBR14.

### Sample JCL to define a GDG model data set using IEFBR14

The following example provides sample JCL to define a GDG model data set, including the DCB information.

```
/* STEP1 - Define the GDG Base
//STEP2   EXEC  PGM=IEFBR14
//BLDDSCB DD  DSN=BPEEXTRC.GDG01,
//          DISP=(NEW,KEEP),
//          UNIT=SYSDA,
//          VOL=SER=PAGE01,
//          SPACE=(TRK,(0)),
//          DCB=(DSORG=PS,LRECL=32756,RECFM=VB,BLKSIZE=32760)
```

You can also pass the DCB information using the DSN parameter on the EXTTRACE statement. However, you cannot override the data set attributes.

You must define the data set with RECFM=VB, and an LRECL that is 4 bytes less than the BLKSIZE. The minimum supported BLKSIZE is 8340.

**Recommendation:** Use a large block size (such as the maximum of 32 760). Using a large BLKSIZE ensures that multiple records are written per block. In addition, you would not have to change your BPE external trace definitions if the minimum supported block size is ever increased.

For more information on IEFBR14, see *z/OS MVS JCL Reference*.

## Starting and stopping BPE external tracing

To start the BPE external trace, you must specify `EXTERNAL=YES` keyword on the `TRCLEV=` parameter of the BPE configuration parameter member (`BPECFG=`) of the `IMS.PROCLIB` data set or issue the `UPDATE TRACETABLE` command with the `EXTERNAL(YES)` parameter on it.

To see which trace tables have external tracing turned on, use the `BPE DISPLAY TRACETABLE` command.

If IMS encounters an I/O error while writing to one of the external trace data sets, IMS closes and deallocates the current data set and allocates and opens a new data set. If IMS cannot allocate or open a new data set, IMS disables BPE tracing for all components that use that GDG data set. Issue the `BPE UPDATE TRACETABLE` command to have IMS retry the allocation and open processing for the data set.

If the I/O capacity of the GDG data sets cannot support the volume of trace records being created, IMS skips writing trace records until the I/O system can catch up rather than waiting and potentially delaying other IMS work.

To stop BPE external tracing for a particular trace table type:

1. Issue the `UPD TRTAB EXTERNAL(NO)` command.
2. Specify the table type on the `NAME` parameter.

To stop BPE external tracing for all tables issue `UPD TRTAB NAME(*) EXTERNAL(NO)`. BPE closes and deallocates the last external trace data set when there are no more table types being externally traced.

Information about BPE commands is in *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*.

### *Formatting and viewing BPE external trace data sets*

To format and view BPE trace entries, use the Interactive Problem Control System (IPCS). For information on using IPCS to view BPE trace entries, see *IMS Version 11 Diagnosis*.

---

## Common Service Layer overview

The *IMS Common Service Layer* (CSL) is a collection of IMS manager address spaces that provide the infrastructure needed for systems management tasks in an IMSplex. The CSL address spaces include Open Database Manager (ODBM), Operations Manager (OM), Resource Manager (RM), and Structured Call Interface (SCI).

The CSL is built on the IMS Base Primitive Environment (BPE) layer. As a result, all BPE externals, such as commands, messages, abends, configurations, and user exit interfaces apply to all CSL manager address spaces.

The IMS CSL provides:

- Improved systems management
- A single system image
- Ease of use through a single point of control
- Shared resources across all IMS systems

A single CSL can service multiple IMS systems within an IMSplex, either on the same operating system or on multiple operating systems.

**Recommendation:** Activate more than one CSL manager of each type in the IMSplex. An IMS control region within an IMSplex that is defined with CSL cannot start unless at least one OM is active in the IMSplex and an SCI resides on every operating system in the IMSplex. This recommendation also applies to RM, although an RM is not required for an IMS control region to start.

There are advantages and disadvantages of having all CSL managers on each operating system. One advantage is better system performance. It is faster for an IMSplex member to communicate with a CSL manager on the same operating system than it is for the IMSplex member to communicate with a CSL manager on a different operating system. A disadvantage is the increased number of address spaces on each operating system. However, if only one instance of a CSL manager is defined in an IMSplex, whether it is ODBM, OM, or RM, no backup of that CSL manager exists if it should fail.

If you do not use shared queues or sysplex technology, you can take advantage of a simplified CSL configuration to issue *type-2 commands* through the CSL OM.

**Related concepts:**

“A simplified CSL configuration” on page 41

## CSL in an IMSplex

The IMS CSL reduces the complexity of managing multiple IMS systems by providing you with a single-image perspective in an *IMSplex*. An IMSplex is one or more IMS subsystems (control, manager, or server) that can work together as a unit.

Typically, but not always, these subsystems:

- Share either databases or resources or message queues (or any combination)
- Run in a z/OS sysplex environment
- Include an IMS CSL

Within an IMSplex, you can manage multiple IMS subsystems as if they were one system. For example, instead of entering commands on each IMS system during local online change, you can enter commands from one single point of control, and the commands run on each IMS system in the IMSplex. An IMSplex can also exist in a local, or non-sysplex, environment. Use of the CSL is optional.

An IMSplex component is an IMS-defined entity that typically runs in its own address space; it manages resources, manages operations, or facilitates communications between other IMS-defined entities. After the necessary started procedures for these components are initialized, they become *IMSplex members*. Examples of IMSplex components are:

- IMS subsystems (DB/DC, DBCTL, DCCTL)
- Resource Manager
- Operations Manager
- Open Database Manager
- Structured Call Interface
- IMS Connect
- A DLIBATCH or DBBBATCH region

A DLIBATCH or DBBBATCH region is considered a special type of IMSplex component in that it does not interact with RMs and OMs.

A typical IMSplex environment including a CSL is shown in the following figure. There are three separate operating system (OS) images. Each OS has an IMS control region and an SCI. In addition, OS 1 has an OM and an RM, OS 2 has an OM and an ODBM, and OS 3 has an ODBM. All three OS images share a coupling facility, which has database sharing structures, a message queue structure, and a resource structure.




Figure 9. IMSplex environment including a CSL

The CSL includes the following IMS address spaces, some of which are optional:

- ODBM
- OM
- RM
- SCI

The address spaces that can participate in an IMSplex are:

- IMS control region address spaces
- IMS CSL manager address spaces (ODBM, OM, RM, SCI)

- IMS server address spaces (Common Queue Server)
- Non-IMS address spaces

The IMSIDs of IMS systems in the IMSplex must be unique.

### **A simplified CSL configuration**

If your IMS configuration does not require RM services, you can still use OM and SCI to take advantage of type-2 IMS commands.

The commands can be issued using the TSO SPOC, the IMS Control Center, or any user-written or vendor-written automation program that uses the OM API.

The IMS Application Menu provides a common interface to enable you to start applications such as TSO SPOC, Syntax Checker, IVP, and more. For more information about the IMS Application Menu, see *IMS Version 11 Installation*.

#### **Related concepts:**

“Common Service Layer overview” on page 38

#### **Related tasks:**

“Defining a simplified IMSplex for the type-2 command environment” on page 25

## **CSL managers**

The CSL is a key component in the evolving architecture of IMS. Using the CSL in an IMSplex provides you with the infrastructure for improved systems management.

The CSL address spaces are also referred to as CSL managers:

- Open Database Manager
- Operations Manager
- Resource Manager
- Structured Call Interface

### **Overview of the CSL Open Database Manager**

Open Database Manager (ODBM) provides distributed and local access to IMS databases that are managed by IMS DB systems configured for either the DBCTL or the DB/TM environments in an IMSplex.

Both independently and with IMS Connect, ODBM supports various interfaces to ease the development of application programs that access IMS databases from many different distributed and local environments. The interfaces that ODBM supports include:

- IMS Universal Database resource adapter
- IMS Universal JDBC driver
- IMS Universal DL/I driver
- The ODBA interface
- The ODBM CSLDMI interface for user-written ODBM client application programs

In addition to supporting the above interfaces, ODBM provides the following functions:



- Supports one-phase commit with or without requiring z/OS Resource Recovery Services as the sync point coordinator. If the ODBM parameter RRS=N is specified, the Database Resource Adapter (DRA) is used to communicate with IMS.
- Routes incoming database requests to IMS systems based on alias names. Alias names for IMS systems are defined to ODBM on the CSLDCxxx PROCLIB member and specified on incoming database requests by the client application programs.  
If the client application program does not specify an alias name on a database access request, ODBM routes the request among multiple IMS systems by using a round-robin method of distribution, in which ODBM routes each incoming request to the active IMS systems that are defined to ODBM.
- Enables ODBM client application programs to access databases from other logical partitions within an IMSplex
- Protects IMS control regions from the unexpected termination during DL/I processing of z/OS application programs that use the ODBA interface through ODBM.
- Functions as an RRMS resource manager capable of participating in RRS controlled sync point processes for ODBA applications that are configured to use ODBM.
- Functions as the resource manager and issues the necessary calls to RRS for single-phase commit processing for local RRS transactions.
- Functions together with IMS Connect as a complete Distributed Relational Database Architecture™ (DRDA®) target server for client application programs that use the DRDA specification.

Because ODBM and IMS Connect support DRDA, you can develop your own DRDA source application server that communicates with IMS by using the distributed data management (DDM) Architecture commands that are a part of the DRDA specification.

ODBM runs in a Common Service Layer (CSL) address space and is an optional IMSplex component. ODBM uses the Structured Call Interface (SCI) services of the CSL for communication and Operations Manager (OM) services of the CSL for command processing.

One ODBM instance must be defined in the IMSplex to use ODBM functions. Each z/OS image can have more than one ODBM. If multiple instances of ODBM are defined in the IMSplex, any ODBM instance can perform work from any z/OS image in the IMSplex.

The following figure shows an overview of an IMS configuration that includes ODBM.




Figure 10. Overview of an IMS configuration that includes ODBM

#### Related concepts:

Chapter 8, "CSL ODBM administration," on page 147



Programming with the IMS support for DRDA (Application Programming)

### Overview of the Operations Manager

OM controls the operations of an IMSplex. OM provides an application programming interface (the OM API) through which commands can be issued and responses received. With a single point of control (SPOC) interface, you can submit commands to OM. The SPOC interfaces include the TSO SPOC, the REXX SPOC API, and the IMS Control Center. You can also write your own application to submit commands.

OM provides the following functions to an IMSplex:

- Routes IMS commands to IMSplex members registered for the command.
- Consolidates command responses from individual IMSplex members into a single response and provides that response to the originator of the command.
- Provides an API for automated operator commands.
- Provides a general use interface to register commands to support any command processing client.
- Provides user exits for command and response edit and command security.

OM log records can be written to a z/OS System Logger log stream by specifying the AUDITLOG= parameter of the IMSPLEX() keyword on the CSLOIxxx PROCLIB member. The log records are an audit trail of command input, associated command output, and unsolicited message output.

One OM must be defined in the IMSplex to use OM functions. Each z/OS image can have more than one OM. If multiple OMs are defined in the IMSplex, any OM defined can perform work from any z/OS image in the IMSplex.

**Related reference:**

“CSL configuration examples” on page 54

## Overview of the Resource Manager

RM helps you manage resources that are shared by multiple IMS systems in an IMSplex. RM provides the infrastructure for managing global resource information and coordinating IMSplex-wide processes, such as global online change.

RM provides the following functions to an IMSplex:

- Maintains global resource information in a *resource structure*, which is a coupling facility list structure that all RMs in the IMSplex can access.
- Ensures resource consistency so that a resource defined as a transaction, lterm, or msname is defined as the same resource type for all IMS systems in the IMSplex.
- Supports resource services.
- Supports client services.
- Uses the Common Queue Server (CQS) to maintain global resource information.
- Coordinates IMSplex-wide processes (such as global online change).

With RM, the system administrator can manage resources that are shared by multiple IMS systems in an IMSplex. RM provides an infrastructure for managing global resources and coordinating processes across the IMSplex.

**Recommendation:** Although a resource structure is optional in an IMSplex, define a resource structure for improved recovery capabilities. However, if you are using shared queues and you want to manage serial applications in an IMSplex, at least one RM must be active with a defined resource structure.

At least one RM must be defined in an IMSplex to use RM functions. You can have one or more RMs on each z/OS image if a resource structure is defined. If no resource structure is defined, you can have only one RM. If you do not require RM functions, you can configure your IMS system without an RM, specifying RMENV=N on the DFSCGxxx PROCLIB member data set. Any RM can process work from any z/OS image within an IMSplex.

**Related reference:**

“CSL configuration examples” on page 54

## Overview of the CSL Structured Call Interface

SCI allows IMSplex members to communicate with one another. Communication between IMSplex members can occur within a single z/OS image or among multiple z/OS images. The individual IMSplex members do not need to know where the other members reside or what communication interface to use.

SCI provides the following functions to an IMSplex:

- Routes messages and requests within an IMSplex.
- Registers and deregisters IMSplex members.
- Notifies IMSplex members when a member joins or leaves the IMSplex.
- Provides security authentication of members when they join the IMSplex.
- Provides a single call interface to isolate the client and server from the underlying communications technology.

Any IMSplex member that requires SCI services must have an SCI on its z/OS image. There can be at most one SCI address space per IMSplex on each z/OS image.

**Note:** DBRC automatic RECON loss notification and parallel RECON access use SCI, but not RM or OM. You can bring up an SCI address space without requiring the other CSL manager address spaces for automatic RECON loss notification.

**Recommendation:** Initialize each CSL manager address space in an IMSplex.

**Related reference:**

“CSL configuration examples” on page 54

## A single point of control (SPOC) program in CSL

A single point of control (SPOC) is a program that allows you to manage operations of all IMS systems within an IMSplex instead of using a master terminal. With a SPOC, you can issue commands to all members of an IMSplex at once.

A SPOC communicates with a single OM. Through SCI, OM then communicates with all of the other IMS control regions in the IMSplex.

A SPOC is optional in an IMSplex. The existing command interfaces for the WTOR, MTO, and E-MCS console are supported only for type-1 commands. If you want to use type-2 commands, you must use a SPOC. You can use a SPOC instead of the z/OS master console, IMS master terminal, or user terminal for most command processing.

Types of SPOCs:

- On a 3270 TSO terminal, a TSO SPOC is a system management application with an ISPF panel interface. You can start the TSO SPOC using the IMS Application Menu. The TSO SPOC also displays the OM audit trail, which has a log of messages issued by IMS and commands issued by operators.
- On a workstation, the SPOC is an IMS system-management application—the IBM Management Console for IMS and DB2 for z/OS (Management Console). The Management Console is an IMS system management application with a graphical user interface (GUI).

- The Batch SPOC utility uses OM services to submit IMS operator commands to an IMSplex. The utility accepts any commands that are supported by the OM API. The Batch SPOC utility is invoked by using standard JCL statements.
- The REXX SPOC API allows automation programs to use SPOC functions. You can use the REXX SPOC API to automate commands.
- Vendor- or user-written SPOC: A program that uses or accesses the OM API to perform SPOC functions.

There can be more than one type of SPOC in an IMSplex, and any number of SPOCs can be active.

A SPOC provides the following functions to an IMSplex:

- Single-system view of an IMSplex by allowing you to operate all IMS systems in the IMSplex from a single console
- Consolidated display of command responses from multiple IMS systems
- Message sending to IMS terminals connected to any IMS control region in the IMSplex

The following figure shows a SPOC application in an IMSplex. The IMSplex in this example has three identical OS images: each has an IMS control region, an IMS CQS address space, and an SCI, OM, and RM. All three OS images share a coupling facility, which includes database sharing structures, a message queue structure, and a resource structure. the IMSplex configuration also includes shared databases and RECON data sets.




Figure 11. SPOC application in an IMSplex

The following figure shows an IMSplex environment with multiple SPOC users. The IMSplex environment includes four IMS control regions, each having its own SCI, and one OM. The multiple SPOC users include two SPOC TSO/ISPF applications, a REXX SPOC API, and a vendor-written SPOC program. Each SPOC can access the IMSplex environment.




Figure 12. Multiple SPOC users in an IMSplex

**Related concepts:**

- ➡ Controlling IMS with the TSO SPOC application (Operations and Automation)
- ➡ REXX SPOC API (System Administration)

**Related tasks:**

- ➡ Starting the IVP from the IMS Application Menu (Installation)

**Related reference:**

- ➡ IMS type-2 command format (Commands)

## REXX SPOC API

The REXX SPOC API allows automated operator application programs written in the REXX programming language to submit IMS operator commands to the IMSplex and to retrieve the command responses.

The command responses are in the form of XML statements. REXX application programs can examine the command responses and perform further processing if needed.

For example, an REXX application program can issue the following command to check the queue count for a transaction:

```
QRY TRAN NAME(prod001) SHOW(QCNT)
```

If the response shows a high queue count, the REXX application program can issue other commands to start additional regions or to change the class of the transaction.

You can write REXX applications programs and run them in a TSO, a batch, or a Tivoli® NetView® for z/OS environment. You must start the SCI, the OM, and the IMS control region address spaces before the REXX application program can issue IMS operator commands.

## Global online change

The global online change function changes resources online for all IMS systems in an IMSplex.

The benefits of global online change are as follows:

- Simplified process for changing resources online by issuing one instance of a command for online change for all IMS systems in the IMSplex
- Continuous processing so that each IMS is not stopped and manually coordinated with online changes.
- No situations where online change is committed on some IMS systems in the IMSplex and not committed on other IMS systems

With global online change, the master IMS control region coordinates online change of the following resources among all the IMS systems in the IMSplex:

- Databases (DMB in ACBLIB)
- Database directories (DDIR in MODBLKS)
- MFS formats (FMTLIB)
- Programs (PSB in ACBLIB)
- Program directories (PDIR in MODBLKS)
- Transactions (SMBs in MODBLKS)
- Routing codes (RCTEs in MODBLKS)

Global online change operates through a master IMS control region, as specified by the user or by OM. OM uses RM, if RM is active, to coordinate the phases of online change with the other IMS systems in the IMSplex. Start online change by issuing the command INITIATE OLC through the OM API. At that time you can specify the master IMS control region. If you do not choose a master IMS control region, OM selects one of the IMS systems as the master IMS control region.

If global online change is specified in the IMSplex, the following components are required:

- An IMSplex defined with CSL and at least one OM
- CQS, if there is a resource structure in the IMSplex

The resource structure is used to save IMSplex-wide process status. It also provides more recovery capability during the online change, should any failures occur.

- OLCSTAT data set, which must be initialized by the Global Online Change utility (DFSUOLC0) before the first IMS cold start

If you exclude RM from the IMSplex by specifying RMENV=N, each IMS system must have its own OLCSTAT data set and that OLCSTAT data set must contain the IMSID of the IMS system that owns it and no other IMSIDs.



- OLC=GLOBAL and OLCSTAT= parameters in the DFSCGxxx PROCLIB member data set
- CSLG= parameter in the IMS EXEC statement

At least one RM and a resource structure are recommended for global online change, but they are not required. Consider the following if you choose not to use RM or a resource structure. If you do not use an RM in your IMSplex, the OLCSTAT data set can contain only the IMSID of the IMS system that owns that OLCSTAT data set. Any attempt to restart an IMS system that contains an OLCSTAT data set with a different or multiple IMSIDs results in an abend. IMS rejects INITIATE OLC and TERMINATE OLC commands that are issued by an IMS system other than the IMS system that owns the OLCSTAT data set.

If you choose not to enable global online change, the MODSTAT data sets in the IMS procedure are optional. The IMS control region sample JCL includes DD statements for the MODSTAT data set and the MODSTAT2 data set (if XRF). If you specify the MODSTAT DD statements for an IMSplex where global online change is enabled, the MODSTAT and MODSTAT2 (if XRF) data sets must also exist.

#### **Related tasks:**

“Defining a simplified IMSplex for the type-2 command environment” on page 25

## **ACB library member online change**

In an IMSplex environment, you can use the ACB member online change (OLC) function to add or change individual members of the ACB library, or the entire ACB library, and bring these new or changed members online without quiescing the IMSplex or refreshing the active ACB library.

Members that are not affected by the member OLC are not quiesced.

You can make online changes to the ACB library in IMSplex environments in which:

- A single IMS system comprises the IMSplex. In this configuration, IMS must use the OLCSTAT data set with a CSL that consists of an OM and SCI. An RM is not required.
- Multiple IMS systems comprise the IMSplex. In this configuration, each IMS system must use the same shared OLCSTAT data set, or each system must use its own local OLCSTAT data set. The CSL must consist of an OM, an SCI, and an RM. Although a resource structure is recommended, it is not required.

Global online change for ACBLIB members can only be performed in an IMSplex.

You can specify whether the ACBLIB member is shared or dedicated in the IMSplex. The following example highlights how the ACBLIB member online change process can be coordinated among all IMS systems that share the OLCSTAT data set:

- OM chooses an IMS system to be the command master IMS.
- The command master IMS coordinates the member OLC process with other IMS systems sharing the OLCSTAT data set using RM.
- ACBSHR= must be specified as Y or N in the DFSCGxxx or DFSDFxxx PROCLIB member.
  - Y– Indicates that all of the IMS systems in the OLCSTAT are using the same active and inactive ACBLIB.

- N– Indicates that each IMS in the OLCSTAT is using its own dedicated active and inactive ACBLIB.
- All IMS systems in the IMSplex that share the OLCSTAT data set must specify the value of ACBSHR=.
- The member OLC updates all the active ACBLIBs, whether or not they are shared.

The ACB member online change uses IMS type-2 commands only. The commands that are involved in performing an ACB library member online change are:

- INITIATE OLC PHASE(PREPARE)
- INITIATE OLC PHASE(COMMIT)
- TERMINATE OLC
- QUERY MEMBER TYPE(IMS)
- QUERY OLC SHOW(RSCLIST)

See “Changing or adding IMS.ACBLIB members online” on page 454 for the steps required to perform an ACB member online change. If you want to fall back to the previous version of the changed resource, perform a full online change process with a full library switch.

## Global TM resource management

Using RM and a resource structure improves your ability to manage Transaction Manager resources in an IMSplex. IMS enforces name uniqueness for LTERMs (VTAM®), nodes (VTAM single-session), user IDs (if the installation requests single-signon enforcement), and users.

IMS also ensures resource type consistency for message destinations and supports global callable services.

### Related tasks:

“Defining and tailoring an IMSplex” on page 23

## Automatic RECON loss notification overview

SCI is the only CSL manager required for automatic RECON loss notification. If you are using the CSL solely for the automatic RECON lost notification, OM and RM are not required. However, SCI must be available on each z/OS image. If the IMS control region is started using the DFSCGxxx or the CSL section of the DFSDFxxx member, OM (and possibly RM) is required.

All DBRC instances using a given RECON must join the same IMSplex. You can have your IMS systems configured so that two or more IMS sharing groups are using a single RECON. This configuration is supported if the sharing groups are part of the same IMSplex. However, if each sharing group is identified as a unique IMSplex, then the RECON must be split into multiple RECONs, one per IMSplex. In addition, if the control region and DBRC address spaces both use SCI, they must also join the same IMSplex. You can also have your IMS systems configured so that two or more IMS sharing groups are part of the same IMSplex, but use unique RECONs.

**Recommendation:** Each RECON in the same IMSplex should have a unique DBRC sharing group ID. This is required if parallel RECON access is enabled for more than one RECON in the IMSplex.

**Related reading:** For detailed information about automatic RECON loss notification configuration, see *IMS Version 11 Database Administration*.

## Configuring an IMSplex with CSL

An IMS control region that is running without CSL must be stopped and restarted to connect to CSL. When CSL is introduced into the IMSplex, the IMS control regions can be stopped and restarted one at a time. The entire IMSplex does not need to be shutdown to introduce CSL.

To configure an IMSplex with CSL, you must initialize at least one of each of the following IMSplex components:

- DB/DC, DCCTL, or DBCTL IMS control region
- DBRC for each IMS control region
- Operations Manager (OM)
- Structured Call Interface (SCI) on each operating system on which an IMSplex member is running
- CQS when message queues are shared or RM is used with the resource structure to manage global resources

It is also strongly recommended that you also initialize a Resource Manager (RM) and a resource structure. Although an RM and a resource structure are not required, you cannot share or manage resources globally.

Optionally, you can have the following in your IMSplex:

- CSL Open Database Manager, which provides distributed and local access to IMS databases that are managed by IMS DB systems configured for either the DBCTL or the DB/TM environments in an IMSplex.
- Data sharing of one or more databases in a DB/DC or DBCTL environment.
- One DL/I address space and 1 to 999 dependent regions for each IMS control region.
- Shared queues.

**Important:** All IMS control regions should share message queues. Otherwise, shared queues should not be used. Running in a mixed environment where some of the queues are shared and others are not can result in an unpredictable or misleading output from type-2 commands, which are entered through the OM API.

- Any number of batch (DLIBATCH or DBBBATCH) regions. OM and RM do not interact with batch regions; therefore, they do not affect management of batch regions.

To include CSL, an IMS control region must be started with the CSLG= parameter, or if the DFSDFxxx member of the IMS PROCLIB data set includes the COMMON\_SERVICE\_LAYER section, the DFSDF= parameter. The minimum configuration for the CSL is to have only one SCI on each operating system where an IMSplex component resides (including IMSplex components like TSO SPOC and CQS), one or more OM's anywhere in the IMSplex, and one or more RM's anywhere in the IMSplex.

How many RM's and OM's to have depends on your requirements. An outline of CSL configuration rules follows:

- Only one SCI can run on an operating system for a given IMSplex. Multiple IMS systems on the same operating system use the same instance of SCI to

communicate within the same IMSplex. If one or more IMS systems on the same operating system are associated with multiple IMSplex systems, you must run a separate instance of SCI for each IMSplex.

- At least one OM must be active in an IMSplex. You can have one or more OM's active on one or more operating systems.
- An RM is not required, but at least one active RM is recommended in an IMSplex. You can have one RM active on any number of operating systems. If you have more than one RM, you must also have a resource structure.

### Configuring the resource structure for an IMSplex

The resource structure, supported by the Common Queue Server (CQS), is the repository for resources managed by RM. CQS manages uniquely named client resources that are shared within a single IMSplex by using the coupling facility list structures.

The resource structure is optional; however, if there is no resource structure certain IMSplex functions cannot be used.

Clients can use the resource structure to share resource information, control block information, and so on. The resource name is unique within the structure. Resources can be updated, deleted, or queried. A primary coupling facility list structure is used to contain the resources. Resource structures do not support an overflow structure.

CQS does not support structure recovery, structure checkpoint, or overflow processing for resource structures. Also, to improve performance, changes to resource structures are not logged. However, for recovery of resources, resource structures can be automatically duplexed by specifying so in the CFRM policy.

CQS defines the attributes of the resource structure the first time it connects to the structure. It defines the resource structure as a persistent, ENTRYKEY list structure and the list entries on the resource structure are defined to contain an adjunct area and a data entry. The resource structure is defined to permit SVC dumps, structure alter, and autorebuild.

**Recommendation:** Use multiple RMs and a resource structure in your IMSplex. Without a resource structure, you cannot use the global resource-sharing capabilities of RM. Additionally, without a resource structure, you can start only one RM. If that one RM goes down, you do not have a backup RM. However, if you do not have a resource structure and are in a DBCTL environment, you can still perform global online changes with one RM.

### Configuring a simplified IMSplex for the type-2 command environment

If you want to use type-2 commands and you have a standalone IMS system or multiple IMS systems that only share databases, you can define an IMSplex without RM. This creates a simplified IMSplex with CSL.

#### Related reading:

- For information on restrictions for online change when an RM is not used, see "Global online change" on page 49.
- For information on defining an IMSplex without an RM, see:
  - "Defining a simplified IMSplex for the type-2 command environment" on page 25
  - *IMS Version 11 System Definition*

- For detailed information about CSL configuration, see “CSL configuration examples.”

## CSL configuration examples

Typically, when an IMS control region (DL/I, DBRC, dependent regions) requires the use of the CSL, SCI, OM, and RM are all required. However, different configurations are possible for the CSL in an IMSplex. For example, you might include ODBM or omit RM in your CSL configuration.

The basic CSL components are the CSL managers, which are:

### Open Database Manager

The Open Database Manager (ODBM) is optional in an IMSplex. You can define additional ODBMs in the IMSplex to enhance performance and availability.

### Operations Manager

At least one OM must be available in the IMSplex. You can define additional OMs in the IMSplex to enhance performance and availability.

### Resource Manager

If any IMS in the IMSplex requires RM services, at least one RM must be available when the IMSplex is initialized. You can define additional RMs in the IMSplex for performance and availability enhancements if a resource structure is used. However, only one RM can be started in an IMSplex if a resource structure is not used. You can configure your IMSplex without an RM and use type-2 commands.

### Structured Call Interface

One SCI is required on each z/OS image where an IMSplex member is running. Only one SCI is allowed on each operating system image for a given IMSplex.

## Sample IMSplex configuration with CSL

The following figure illustrates a sample IMSplex configuration that includes the CSL, a SPOC, and automated procedures.

- The OS image includes address spaces for ODBM, OM, SCI, RM, an IMS control region, IMS Connect, and IMS CQS.
- The OS image shares a coupling facility and databases.
- A SPOC application, an automation application, a master terminal, and an end user terminal all access the OS image.
- The Open database resource adapters and APIs are also shown accessing the OS image through IMS Connect, the TCP/IP gateway for IMS.




Figure 13. Sample IMSplex configuration with CSL

To obtain the fastest communication, define an OM and RM on each z/OS image; an IMSplex component can communicate more quickly with an OM or RM on the same z/OS image as that component rather than an OM or RM on a different z/OS image. However, this configuration also increases the number of address spaces on each z/OS image, which can create a more complex operating environment.

**Note:** If only one RM and only one OM are defined across an IMSplex, there is no backup to perform that manager's work in case of failure.

**Recommendation:** Define more than one RM, OM, and SCI across an IMSplex.

## IMSplex minimum CSL configuration

The following figure illustrates the minimum configuration possible for the CSL in an IMSplex. Each OS image has an IMS control region and an SCI; in addition, the first OS image also has an OM, but no RM. RMENV=NO was specified on the startup parameter.




Figure 14. IMSplex minimum CSL configuration

In the figure, each z/OS image has a separate SCI since each has a distinct IMS control region. OM can reside on one z/OS image and still be used by other images in the IMSplex.

## IMSplex mixed version CSL configuration

The following figure illustrates a more complex configuration having multiple versions of IMS within the IMSplex.

- Operating System 1 has IMS Version 12 control regions and is running OM and RM
- Operating System 2 has IMS Version 13 control regions, an ODBM, and an OM
- Operating System 3 has an IMS Version 11 control region

All three OS images share the coupling facility. Within the coupling facility is a resource structure. Therefore, an additional RM could be defined in another OS image.




Figure 15. IMSplex mixed version CSL configuration

## IMSplex DBCTL CSL configuration

The following figure illustrates the configuration of the Common Service Layer in a DBCTL environment. Here, three OS images each have an IMS DBCTL control region. All have an SCI and an OM. Two have an ODBM. Only one RM is allowed in the IMSplex, here in OS1, because no resource structure is defined.

**Recommendation:** Define a resource structure for DBCTL.






Figure 16. IMSplex DBCTL CSL configuration

In the figure, an OM is defined in each OS image to enhance overall system performance. If you do not use global online change, neither RM services nor an RM address space is required. In such a configuration, the CSL includes SCI and OM.

## IMSplex minimum ODBM configuration

The following figure illustrates a minimum configuration for ODBM in an IMSplex that include three IMS systems. Each OS image has an IMS control region, an SCI, and an ODBM; in addition, the first OS image also has an OM, which is required by ODBM, but no RM. RMENV=NO was specified on the startup parameter.

The ODBM configuration shown is not a data sharing environment. Each IMS system manages its own databases separately.

Also, the databases on operating system 3 cannot be accessed by ODBM clients, such as the IMS Universal drivers, because an instance of ODBM is not running on operating system 3.




Figure 17. IMSplex minimum ODBM configuration

In the figure, each z/OS image has a separate SCI since each has a distinct IMS control region. OM can reside on one z/OS image and still be used by other images in the IMSplex. An instance of ODBM must be active on each z/OS image that contains databases to which ODBM clients require access.

### Shared queues in an IMSplex without a CSL

In the following figure, three OS images are each defined with an IMS control region and IMS CQS. Each is associated with a Master Terminal Operator (MTO) console. All three OS images share a coupling facility that includes database sharing structures and a message queue structure. No CSL manager address spaces are defined.

This figure illustrate what a shared queues IMSplex environment looks like without a CSL.




Figure 18. Shared queues in an IMSplex without a CSL

### Shared queues in an IMSplex environment with a CSL

In the following figure, three OS images are defined. Each is identical, having an IMS control region, IMS CQS, and all CSL managers (SCI, OM, and RM). The three OS images share the coupling facility, which includes database sharing structures, a message queue structure, and a resource structure. A SPOC application can access any of the OS images.

This figure illustrates what a shared queues IMSplex environment looks like with a CSL.




Figure 19. Shared queues in an IMSplex environment with a CSL

### IMSplex single system CSL configuration

In a single system IMSplex environment, each address space of the Common Service Layer is defined in a single OS image. This simple configuration is shown in the following figure. The OS image is defined with multiple IMS control regions, and one SCI, RM, OM, and ODBM.




Figure 20. IMSplex single system CSL configuration

**Related concepts:**

“Overview of the Operations Manager” on page 43

“Overview of the Resource Manager” on page 44

“Overview of the CSL Structured Call Interface” on page 45

**Related tasks:**

“Defining a simplified IMSplex for the type-2 command environment” on page 25

## Common Queue Server overview

Common Queue Server (CQS) is a generalized server that manages data objects on a *coupling facility list structure*. CQS manages two types of structures: *queue structures* and *resource structures*. CQS receives, maintains, and distributes data objects from these structures on behalf of clients. IMS is one example of a CQS client that uses CQS to manage both its shared queues and shared resources.

**This topic contains General-use Programming Interface information.**

CQS uses the z/OS coupling facility as a repository for data objects. Storage in a coupling facility is divided into distinct objects called *structures*. Authorized programs use structures to implement data sharing and high-speed serialization. The coupling facility stores and arranges the data according to list structures. Queue structures contain collections of data objects that share the same name, which are known as *queues*. Resource structures contain data objects that are organized as uniquely named resources.

CQS runs on z/OS. The CQS client must also run on z/OS. CQS runs in a separate address space that can be started by the client.

Clients communicate with CQS by using CQS requests that are supported by CQS macro statements. Using these macros, CQS clients can communicate with CQS and manipulate client data on shared coupling facility structures. The following figure shows the communications and the relationship between clients, CQSs, and

the coupling facility.




Figure 21. Client systems, CQS, and the coupling facility

#### **Related reading:**

- CQS requests are described in *IMS Version 11 System Programming APIs*.
- See *z/OS MVS Setting Up a Sysplex* for complete details about setting up a sysplex.

### **CQS benefits**

CQS enables users to take advantage of the benefits of a Parallel Sysplex (sysplex) environment. These benefits include:

#### **Automatic work load balancing**

CQS places data objects on shared queues where they can be processed by any participating client system.

Any participating client system can use CQS to retrieve a data object from the shared queues.

#### **Incremental growth**

Customers can add new systems as workload increases.

#### **Reliability**

For both shared queues and resources, if one client system fails, the remaining client systems process the work.

### **CQS components**

CQS uses the following components:

#### **Primary structure**

A z/OS coupling facility list structure that contains shared queues.

**Resource structure**

A z/OS coupling facility list structure that contains uniquely named resources.

**Overflow structure**

A z/OS coupling facility list structure that contains shared queues when the primary structure reaches an installation-specified overflow threshold. The overflow structure is optional.

**z/OS log stream**

A shared z/OS log stream that contains all CQS log records from all CQSs connected to a structure pair. This log stream is important for recovery of shared queues, if necessary. Each structure pair has an associated log stream.

**Checkpoint data set**

A local data set that contains CQS system checkpoint information.

**Structure recovery data sets (SRDSs)**

Shared data sets that contain structure checkpoint information for shared queues on a structure pair. Each structure pair has two associated SRDSs.

**CQS functions**

CQS performs the following functions:

**Processes CQS requests**

An architected interface that clients use to access CQS or data objects on a queue structure or a resource structure.

**Notification of work on a queue**

Clients register interest in the shared queues. If an empty queue becomes non-empty, CQS notifies its registered clients.

**Records restart and recovery information**

CQS records all the information necessary for restart and recovery in the z/OS system logger.

**CQS system checkpoint**

CQS system checkpoint writes log records relating to a particular CQS to the CQS log. The log records contain information necessary for CQS to restart and recover work.

**Structure checkpoint**

The structure checkpoint copies the queues from a structure pair into an SRDS for recovery purposes.

**Structure rebuild**

Structure rebuild is a z/OS process that allows another instance of a structure to be allocated with the same name and data reconstructed from the initial structure instance.

**Overflow processing**

CQS provides an overflow option to help prevent a queue full condition. When the primary list structure reaches the overflow threshold value, CQS attempts to dynamically increase the size of the primary structure, offloads selected queues to an overflow structure, or rejects requests for selected queues.

## CQS requirements

The CQS has coupling facility requirements that are documented in the *IMS Version 11 Release Planning*.

**Related reading:** See *z/OS MVS Setting Up a Sysplex* for more information on:

- The planning required to migrate to a sysplex that uses a coupling facility
- Hardware configurations of a sysplex
- The software products that can use the coupling facility
- The tasks for migrating to a coupling environment
- Checklists for installing the sysplex hardware and software

## Queue structures

A queue structure is a coupling facility list structure that contains a collection of data objects, some of which might have the same name. Data objects that have the same name are considered to be on the same queue.

Queue structures support structure overflow, in which an associated overflow structure can be allocated to prevent the queue structure from becoming full. A primary queue structure and its associated overflow structure are known as a structure pair.

CQS physically divides the queue structure list headers into 11 private queue types for CQS use and 11 client queue types for client use. Client queue types are defined by the client. A client can group queues associated with a type of work, such as transactions. A queue type can have a value of 1 to 255. Any queue type over 11 is mapped into one of the physical queue types.

CQS manages private queues and client queues on queue structures. CQS uses the private queue types to manipulate client data objects for CQS requests. Each client queue type can be used by a client for a different type of work. A client registers interest in only those queue types that it can process, based on the types of work you define for it.

### Private queue types managed by CQS

The following table shows five of the private queue types, and the work that a client processes on them.

| Queue type    | Description                                                                                                                                       |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Cold queue    | Contains data objects that are in doubt for a client or for a CQS that cold started                                                               |
| Control queue | Contains control list entries that CQS uses to manage list structures and control processes (such as structure checkpoint and structure recovery) |
| Delete queue  | Intermediate queue used for CQSDel request processing                                                                                             |
| Lock queue    | Contains data objects that are locked by the CQSREAD request                                                                                      |
| Move queue    | Intermediate queue used for CQSMOVE request processing                                                                                            |



## Resource structures

A *resource structure* is a coupling facility list structure, used by the Common Service Layer's Resource Manager and managed by CQS, that contains uniquely named resources.

This structure is typically used to maintain global resource information when multiple Resource Managers exist in an IMSplex. Resource structures enable CQS to perform resource management in an IMSplex.

CQS physically divides the resource structure list headers into 11 private resource types for CQS use and 11 client resource types for client use. Client resource types are defined by the client. A resource type can have a value of 1 to 255. Any resource type over 11 is mapped into one of the physical resource types.

Clients can use the resource structure to share resource information, control block information, and other types of information. The resource name is unique within the structure. Resources can be updated, queried, or deleted. A primary coupling facility list structure is used to contain the resources.

## CQS structure functions

CQS provides functions for monitoring structure status and capacity, and enabling structure recovery. Some of these functions are built-in functions and do not require intervention. Other functions are optional, and can be set up or initiated as your installation needs them.

### Structure overflow

CQS provides a structure overflow function that automatically warns you when a queue structure is approaching full and takes action to prevent a full structure. When the usage of a structure reaches the overflow threshold, CQS attempts to make the structure larger by initiating a structure alter. If the alter fails, CQS either allocates an overflow structure and moves selected queues to the overflow structure (if you define an overflow structure), or prevents new data objects from being put on the selected queues.

**Important:** Overflow processing is not supported for resource structures.

### Structure rebuild

Structure rebuild is a z/OS process that allows another instance of a structure to be allocated with the same name and contain data reconstructed from the initial structure instance. z/OS supports system-managed rebuild, in which case z/OS rebuilds the structure. z/OS also supports user-managed rebuild; the user rebuilds the structure. Structure rebuild can be initiated manually by using an operator command, or automatically by CQS or z/OS.

CQS allows system-managed rebuild for queue structures and resource structures. CQS provides user-managed rebuild to support a structure copy function and a structure recovery function.

Structure copy copies the contents of a structure to another structure for planned reconfiguration. Structure copy is supported for resource structures and queue structures.

Structure recovery recovers a structure from the structure checkpoint data set and the CQS log after a structure failure. Structure recovery is supported for queue structures.

## Structure duplexing

CQS can use the *duplexing* capabilities of z/OS. Duplexing occurs when the operating system creates a duplex (backup) copy of a structure, then maintains the two structures during normal mainline operation. If a structure fails, or a connection to a structure is lost, the operating system switches to the unaffected structure instance.

### Related concepts:

“Preventing CQS structure full” on page 182

“Rebuilding structures in CQS” on page 185

## CQS recovery functions

CQS provides functions for recovering work-in-progress, queues, and resources in case of system shutdown or failure. Some of these recovery functions are built-in and do not require intervention.

Other functions are optional and can be set up or initiated as you need them.

## System checkpoint

To enable CQS restart in the event of failure, CQS periodically takes a “snapshot” of all control blocks and tables, and writes that information to the z/OS log. That process is called *system checkpoint*. System checkpoint can be initiated by CQS, the client, or manually with an IMS command.

## CQS logging and the z/OS system logger

CQS uses the z/OS system logger to record information necessary for CQS to recover queue structures and restart. CQS writes log records for each recoverable coupling facility list structure pair that it uses to a separate log stream. The log stream is shared among all CQS address spaces that share the structure. The system logger provides a merged log for all CQS address spaces that are sharing queues on a coupling facility list structure.

**Important:** Changes to resource structures and nonrecoverable (RECOVERABLE=NO) queue structures are not logged.

## Structure checkpoint

To enable queue structure recovery in case of failure, CQS periodically takes a “snapshot” of the queues on all queue structures. That process is called *structure checkpoint*. Structure checkpoint can be initiated by CQS, the client, or manually with an IMS command.

**Important:** Structure checkpoint is not supported for resource structures.

## CQS client requests

CQS client systems communicate with CQS using a general use interface consisting of CQS requests.

**Related concepts:**

“Recording information necessary for starting CQS” on page 177

“Using CQS structure checkpoint” on page 181

 CQS client requests (System Programming APIs)

---

## Type-2 command environment

The type-2 command environment requires a Common Service Layer (CSL) with Operations Manager (OM) and the Structured Call Interface (SCI). A Resource Manager (RM) is not required. Any IMSplex (including a single-IMS IMSplex) with a CSL constitutes a type-2 command environment.

The type-2 command environment can include a single, standalone IMS system, multiple IMS systems that do not share resources, or a full-feature IMSplex that includes RM. The type-2 command environment can also be used with DB/DC, DCCTL, and DBCTL environments.

Some type-2 commands (such as the type-2 commands associated with the Open Database function) have additional requirements.

With a type-2 command environment, you can also use OM command security. You might also need to consider SCI security in a type-2 command environment.

Type-2 commands must be issued from an automated operator application program, such as the TSO single point of control (SPOC) application that is shipped with IMS. IMS offers two interfaces that such application programs can use to communicate with IMS: the Operations Manager (OM) application programming interface (API) and the REXX SPOC API.

**Related concepts:**

 IMSplex overview (System Administration)

 Controlling IMS with the TSO SPOC application (Operations and Automation)

 REXX SPOC API environment with the CSL OM (System Programming APIs)

**Related tasks:**

 Defining a simplified IMSplex for the type-2 command environment (System Administration)

**Related reference:**

 IMS type-2 command format (Commands)

---

## Extended Terminal Option

The Extended Terminal Option (ETO) is a function of IMS TM that can be included at system definition.

With ETO:

- You can add VTAM terminals to IMS without redefining the system.
- You can dynamically add user LTERMs and remote LTERMs (for MSC links) to IMS.
- You cannot define the master terminal using ETO.
- You cannot use the XRF surveillance link with ETO.

#### **Related reading:**

- For more information on specifying ETO at system definition, see *IMS Version 11 System Definition*.
- For information on ETO system security features, see Chapter 18, “IMS security,” on page 307.

---

## **APPC**

IMS supports advanced program-to-program communication (APPC) conversations in two scenarios: APPC/IMS and the explicit CPI-C driven interface.

The two scenarios differ in the subsystem that manages the updating and synchronization of protected resources that the application program accesses.

In the APPC/IMS scenario, when SYNCLVL = NONE or SYNCLVL = CONFIRM, IMS is the synchronization-point manager. When SYNCLVL = SYNCPT, z/OS Resource Recovery Services (RRS) is the synchronization-point manager.

In the CPI-C driven scenario:

- The sync point manager is the RRS function of z/OS.
- The resource manager is IMS.
- The program that accesses and updates the protected resources is the APPC/MVS application program.

APPC/IMS defines the formats and protocols for program-to-program communication. APPC/IMS enables applications to be distributed throughout the network and to communicate with each other regardless of the underlying hardware architectures and software environments. APPC/IMS provides a facility for implementing logical unit type 6.2 (LU 6.2) support.

APPC/MVS is used for all interaction with remote LU 6.2 devices or subsystems. IMS accesses the session through APPC/MVS services. Using APPC/IMS, IMS and LU 6.2 devices access each other without requiring coding changes to existing application programs. With slight modifications to IMS application programs, Common Programming Interface (CPI) communications-driven application programs can communicate with IMS application programs (and can execute as IMS application programs).

A restriction exists, however, on LU 6.2 synchronous conversations with implicit transactions. If a transaction creates more than one daughter transaction, which, in turn, might create other transactions, and one of the daughter transactions provides the response, then the result is unpredictable. In some cases, depending on the execution sequence of the transactions, the LU receives a DFS2082 message and the response is sent to the default TP name DFSASYNCR. In other cases, the LU receives the response and no DFS2082 message is issued.

#### **Related information:**

 3303 (Messages and Codes)

---

## **Security for dependent region processing**

Although security checking can be carried out by terminal, transaction, command, and other types of authorization, you can also implement security by limiting the resources that application programs that are scheduled in dependent regions can access using resource access security (RAS).

RAS uses RACF®, RACF security classes, and user IDs to define resources and the dependent regions that can use those resources. To implement RAS security, you must define in the RACF security classes resource profiles for the transactions, PSBs, and LTERMs that you want to protect. You must also specify in the resource profiles the user IDs of each dependent region that you want to allow to use each resource.

RACF is an external security product to IMS, accessed by IMS using the Security Access Facility (SAF). RACF is licensed with the IBM z/OS Security Server. Where this information directs you to use RACF, you can use a different, equivalent security product if you choose.

When an application program executing in a dependent region attempts to access a resource, RACF checks the resource's security class profile to see if the user ID of the dependent region in which the application resides is authorized for that resource. If the resource profile lists that user ID, RACF allows access; if not, RACF denies access.

You can also use exit routines, such as the Resource Access Security exit routine (DFSRAS00) with RAS, which allows you to customize the security checking for dependent region processing.

---

## MPP scheduling in DB/DC and DCCTL environments

After an MPP region has been initialized, it can execute an application program within its virtual storage. By using the scheduling algorithm, the control program selects a message for processing. Using the transaction code, the appropriate application program is loaded into dependent region storage from the IMS.PGMLIB data set.

The application program is identified by the PSB name that was declared in system definition to be associated with that transaction code. The convention used by IMS TM for MPPs is that the application program name is the same as the PSB name. The first message segment is then made available from the message queue, and control is passed to the application program.

The scheduling algorithm also controls the amount of processing performed by the application program. You can specify a limit to the number of messages processed in the scheduling of one program. When this number is reached, IMS TM does the following:

- If any equal or higher priority transactions are queued, IMS TM terminates the application program. The region becomes available for another program to be scheduled into its storage. IMS uses the scheduling algorithm to choose the program to schedule.
- If no equal or higher priority transactions are queued and messages are still queued for the current application program, the region goes through quick reschedule and returns the next message to the application program.
- If equal priority transactions are enqueued, IMS allows the transaction to quick reschedule if the PROCLIM value has not been reached. If the PROCLIM value has been reached, IMS disallows the quick reschedule and processes the other transactions.

Quick rescheduling is also affected by the IMS scheduling algorithm, which takes into account the following factors and more:

- MAXRGN value

- PARLIM value
- Current number of regions in which the transaction is scheduled
- If no more messages exist for the scheduled transaction, IMS TM determines if other work for the region is ready to be processed.
- If no additional work is ready to be processed, IMS TM determines if the region can become pseudo wait-for-input (pseudo WFI). This determination causes one of the following actions:
  - If the region is eligible for pseudo WFI, the region remains scheduled for the transaction and waits until another message is entered for the region. If the next message is for the scheduled transaction, the message is passed to the application program. If the next message is for a different transaction, IMS TM terminates the application program and schedules a new application program to process the new message.
  - If the region is not eligible to become a pseudo WFI, IMS TM tells the application program that no more messages exist, and the application program terminates.

The master terminal operator directly schedules batch message programs by the entry of the JCL to start the batch message region. The program and PSB to be used are explicitly given in the EXEC statement.

---

## Scheduling application programs against unavailable data

IMS Transaction Manager schedules an application program even when some of the full-function databases that the application program can access are not available. When dealing with unavailable data, the application program can be sensitive or insensitive to data unavailability.

Application programs that are sensitive to unavailable data issue an INIT call when IMS schedules the applications. The INIT call tells IMS to return a status code in the PCB when the data that the application program requires is not available. The program can then take the appropriate action.

Application programs that are insensitive to unavailable data do not issue an INIT call. If an application program that is insensitive to unavailable data attempts to access data that is not available, IMS terminates the application program with a user abend code 3303 and backs out any updates that the application program has made. If an application program generates ten 3303 abends, IMS prevents further rescheduling of the application program by stopping its PSB.

If an application program terminates as a result of unavailable data, IMS places the input message that the application program was processing on the suspend queue. A separate suspend queue exists for each transaction type. Serial transactions are not placed on the suspend queue, but rather are returned to the normal queue as the next message to be processed and are USTOPPED.

If IMS determines that most messages are failing and being placed on the suspend queue, IMS stops processing that transaction type. When the transaction is started, or when a database used in processing the transaction is started, the messages on the appropriate suspend queues are transferred to the normal queue, and another attempt is made to process the message.

### Related reading:

- For additional information about scheduling application programs against unavailable data, see *IMS Version 11 Application Programming*.
- For more information on the suspend queue, see “Scheduling transactions using the suspend queue” in *IMS Version 11 System Definition*.

---

## Fast Path

Use Fast Path to improve performance for simple transactions. When data communication requirements are for a high transaction volume with rapid database updates and inquiries, the Fast Path facilities offer several advantages over full-function DL/I processing.

Examples of application programs with these requirements are the teller transactions in banking and point-of-sale transactions (inventory update) in retail marketing. Fast Path input and output messages use the expedited message handler (EMH), bypass message queuing, and priority scheduling. Most terminals have Fast Path execution potential. However, terminals that cannot run in response mode do not have Fast Path potential.

For a DB/DC environment, Fast Path requires the Database Manager and Transaction Manager and becomes an integral part of the IMS online system. The control program manages concurrent processing of Fast Path and DL/I programs.

The DCCTL environment supports Fast Path processing and transactions, but not Fast Path databases.

### **Related reading:**

- For more information on the design, definition, initialization, monitoring, or tuning of databases used with Fast Path, see *IMS Version 11 Database Administration*.
- For information on Fast Path application programming, see *IMS Version 11 Application Programming*.

## Fast Path databases

In addition to DL/I databases, two other database types are available with Fast Path: the main storage database (MSDB) and the data entry database (DEDB). These databases are designed for the kinds of application programs that require high availability. The two types offer a choice of either rapid response within high activity or partitioned access within a large volume of data.

**Related reading:** For a detailed description of the design advantages and implementation of these Fast Path databases, see *IMS Version 11 Database Administration*.

## Dependent region use for Fast Path

The majority of Fast Path processing programs are similar in function to message processing programs (MPPs). Message-driven programs correspond to MPPs, and execute in a Fast Path dependent region (IFP region). These programs execute in wait-for-input mode, so that the program execution is equivalent to a dependent region operation. Parallel scheduling is supported, so that another copy of the program can execute in another dependent region.



Because of the ability to perform much of the data entry database maintenance online, such as reorganization and recovery-related functions, your IMS online system should allow for the scheduling of a Fast Path utility region.

Fast Path application programs and utilities can be active concurrently with message processing programs or BMPs. An IMS online system that is using Multiple Systems Coupling (MSC) can also be processing Fast Path transactions. However, message input received through MSC links cannot be directed to a Fast Path application program, nor can they be passed to the Fast Path input exit routine. This restriction does not apply to message input received using Intersystem Communication (ISC) connections.

## **Fast Path transactions**

A Fast Path application program is driven by transactions that bypass IMS input message queue handling. A transaction can be declared to be Fast Path exclusive. After initial edit, the input message is passed to an exit routine. This routine helps determine the dependent region in which the transaction is executed. The message is added to a Fast Path message-handling area in the program's storage, and then the transaction is made available to the message-driven program without I/O to the message queues.

A transaction can also be declared as having Fast Path potential. After entry, the transaction is also passed to the user exit routine, which decides whether the transaction should pass directly to the message-holding area in the control program's storage, or whether it should be routed to IMS for normal message queue handling. The queue bypass again leads to the transaction being presented to a message-driven program.

The control of Fast Path messages within the storage of a control program is called expedited message handling. One of the checks ensures that messages meet the requirement that single-segment input and output messages are used. The Fast Path Input Edit/Routing exit routine is DBFHAGU0. IMS can use EMH buffers for complete input editing of both Fast Path and full-function transactions.

## **Fast Path in a DBCTL environment**

In a DBCTL environment, Fast Path provides improved performance and data availability to programs that can use data entry databases (DEDBs). A DBCTL environment supports only the Fast Path facilities that are related to DEDBs. You cannot run main storage database (MSDB) facilities.

BMPs or CCTL threads can schedule a PSB to access DEDBs. Parallel scheduling is supported; another copy of the PSB can execute in another BMP or CCTL thread. Fast Path application programs and utilities can be active concurrently with BMPs.

Because much of the DEDB maintenance (such as reorganization and recovery-related functions) can be performed online, your IMS DBCTL environment should allow for the scheduling of a Fast Path utility region.

---

## **Automated operator application programs**

An application program that can issue a subset of IMS operator commands using the DL/I CMD or ICMD calls is called an automated operator (AO) application program.



AO application programs can use the CMD call in only the DB/DC and DCCTL environments.

AO application programs can use the ICMD call in all environments.

When the CMD or ICMD call is issued, the operator command is executed and the first segment of the command response is put in the AO application program's I/O area. Subsequent parts of the response are obtained by GCMD (if CMD is issued) or RCMD (if ICMD is issued) calls.

To maintain security, you need to decide which AO application programs can issue operator commands and which commands they can issue. An AO application program can issue a single command or a series of commands.

**Related reading:**

- For information on securing the CMD and ICMD calls, see Chapter 18, “IMS security,” on page 307.
- For more information on automated operator (AO) application programs, see *IMS Version 11 Operations and Automation*.

When using an AO application program it is also possible to automate the recovery of databases and areas by using the IMS AO facility in conjunction with a database recovery service. One way that you can control recovery is by using the ICMD call to issue /RECOVER commands in a type 2 AO application program.

**Related reading:** For more information on the /RECOVER commands, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

---

## System logging and processing continuity

To protect the integrity of the data, the online IMS uses both external security checking and various internal techniques to record the transactions entered into the system and the database update activity.

The principal tool for recording online system activity is IMS system logging. Data stored on various logging data sets contains information used for restart, recovery, statistics, and audit purposes.

IMS log data is recorded in four kinds of data sets:

- Online log data set (OLDS)
- Write ahead data set (WADS)
- Restart data set (RDS)
- System log data set (SLDS)

The online system uses a minimum of three OLDSs, one WADS, and a single RDS, all residing only on DASD. When one or more online log data sets are filled, you can archive them to system log data sets using the IMS Log Archive utility. DASD or tape media can be used for SLDSs. Batch systems use system log data sets and are able to log to either tape or DASD.

The online system uses the OLDSs in wrap-around fashion. If dual logging of the OLDSs is an installation requirement, a pair of data sets (primary and secondary) must be assigned. The DD names for the OLDSs begin with the character string DFSOLP for the primary data set and DFSOLS for the secondary data set. A

unique suffix (00 through 99), called an “OLDS identifier,” completes the 8-character DD name. Either single logging or dual logging is performed, as determined by DD statements during system initialization or by instructions included in the DFSVSMxx IMS.PROCLIB member.

A WADS is a small data set containing a copy of log records that are in OLDS buffers but have not yet been written to the OLDSs. When logging to DASD (required for online processing), fixed-length blocks make direct retrieval easier. A WADS allows large fixed-length blocks (in variable blocked format) to be written to the OLDSs without the requirement to rewrite blocks. When log data has been written to the OLDSs, the WADS is reused.

If a system failure occurs, the log data in the WADS is used to close the OLDSs. The close process occurs as part of an emergency restart or as an option of the Log Recovery utility.

Write-ahead support is provided for a spare WADS. When a write error is detected, a spare WADS replaces the WADS that encountered the error. Dual WADS logging is also supported if it is required to have backup in the event of a read error while closing the OLDSs from the WADS.

The online IMS system controls the log data sets that are used for startup. It makes use of entries in the checkpoint identification table written on the restart data set, and of log data set information recorded in the DBRC RECON data set. If you are using automatic restart, the /START IMS command issued from the system console causes the appropriate kind of restart. Normally, this restart results in the use of OLDSs records and in a normal restart that completes without the use of prior system logs. If restart processing abnormally terminates before the initial checkpoint, the appropriate restart for automatic restart is the same type (/NRESTART or /ERESTART) as the aborted restart.

---

## Checkpointing

Checkpointing is the primary technique that IMS uses to record information that can be used to restart an interrupted operation. Using the status information recorded during a checkpoint, IMS restores the contents of the message queues and database changes.

Checkpoints are an integral part of system shutdown and startup. Also, the amount of reprocessing, back from the point of system interruption and forward to a continuation point, is reduced when checkpointing is reasonably frequent. Some processing overhead is associated with checkpoint information, but this is an acceptable trade-off for the efficient restart of the system.

In an XRF complex, SNAPQ checkpoint records taken on the active IMS system are used to build control blocks on the alternate IMS system during the synchronization phase.

IMS internal checkpoints are scheduled to occur automatically at predetermined intervals. The interval is specified in terms of an increment to the number of system log records created. As the online IMS events are logged with individual log record types, a count is maintained. When the increment exceeds the specified value, checkpoint processing is invoked. IMS checkpoints can also be invoked explicitly by the master terminal operator and by application programs that have been authorized to issue the /CHECKPOINT command.

Fast DB Recovery regions monitor checkpoint records on IMS systems, and uses them during database resource recovery.

---

## Locking mechanisms and database integrity in DB/DC and DCCTL

IMS offers a choice of locking; you can use program isolation (PI) locking or the services of the Internal Resource Lock Manager (IRLM). The IRLM component is used as an integral part of data sharing.

With program isolation, all activity (modifying the database and creating messages) of an application program that is active in the DB/DC environment is isolated from any other application programs that are active in the system. The isolation persists until that application program confirms, by reaching a synchronization point, that the data it has modified or created is valid.

**Note:** A limitation of the PI lock manager is that it can only manage up to 63 applications that are waiting to own a lock. If one additional application asks the PI lock manager for a lock (this application would be waiter number 64), the application terminates with an abend 2478, is backed out by IMS, and its locks are released. If the requestor is an IMS TM transaction, it is returned to the queue and is reprocessed.

The locking mechanisms are also used to:

- Remove the effects of an abnormally terminated application program
- Perform the processing required for a ROLL, ROLB, or ROLS call
- Resolve deadlock situations

For all the above processing, the removal of database updates and held output messages is done from the previous synchronization point up to the current status. A *synchronization point* is defined as the point at which an application program can be restarted. The first such point for an application program is its initial scheduling. The most common synchronization point is when a GU to the message queue occurs. By issuing a call for the next message, a program in single message mode is indicating the start of a cycle of processing and the completion of any previous work. At this time, any output messages that are queued to a temporary destination are sent to their final destination, and database updates are committed.

An application program can also issue a CHKP call, which forces a synchronization point. For application programs executing in multiple message mode, or BMPs that are not transaction driven, the synchronization point is the time of either the initial scheduling or the last CHKP call.

Another aspect of program isolation is the control of database updates at the segment occurrence level. During the scheduling process, IMS analyzes the intent of an application program toward the database it uses. If a conflict exists with the database usage of a currently scheduled transaction and a candidate for scheduling because an application program needs exclusive use of the database, the scheduling process must select another transaction code and try again. If exclusive intent is not a factor (this is usually the case), application programs are scheduled concurrently. IMS controls the interleaved ownership of database segments with a locking mechanism. As application programs execute, they enqueue on the database records and release those resources, either after update or when the application program reaches a synchronization point.

Possible deadlock situations are resolved in a manner transparent to application programs and terminal operators. When IMS detects a deadlock situation, one of the application programs involved in the deadlock is abnormally terminated with a special abnormal termination code. The abnormal termination causes the activity of the terminated program to be dynamically backed out to a previous synchronization point. Its held resources are released. This allows other application programs to complete their processing. The special code causes the transaction that was being processed to be saved. The application program is rescheduled.

In DBCTL, if a deadlock situation forces an abnormal termination of a CCTL thread, that thread is not saved or retried by DBCTL. The CCTL, upon receiving certain deadlock termination codes, retries its transaction.

If a BMP is selected to be dynamically backed out, it cannot be rescheduled and terminates at its latest synchronization point. If the BMP did not access the message queues for input or issue CHKP calls, the BMP terminates during scheduling and all the BMP database update activity is nullified.

If an ODBA thread is active when IMS DB shuts down or terminates abnormally, the ODBA application thread is terminated. The ODBA application program is not terminated, but is no longer able to make calls on the thread.

---

## Exit points and routines

You can modify how IMS processes a unit of work by providing IMS with subroutines (called *exit routines*) that perform special processing. IMS calls exit routines at various points in its logic flow (called *exit points*) and allows you to control how IMS performs its work.

One example of modifying how IMS processes a unit of work might be to interrogate an in-flight message (in an MSC environment) to determine if it originated from a test or production system. If it came from a test system, you might want to route that message to an application that collects data about the test system. If the message originated from a production system, you probably would not want to change its destination.

To perform the interrogation of the message, you would code a TM and MSC Message Routing and Control User Exit routine. If you name this routine DFSMSCE0, place it in the IMS.SDFSRESL library, and bind it to that library (or a concatenated library), IMS calls your exit routine when IMS receives the message. When your exit routine is done with its processing, the exit routine returns control to IMS and then IMS resumes processing the message to either the original destination or the test application.

Certain exit routines are required and others are optional. Some IBM-supplied exit routines can be used as is and some require modification before using.

**Related reading:** For complete information about exit points and exit routines, see *IMS Version 11 Exit Routines*.

## Data Capture exit routines

If your installation contains both IMS DB and DB2 for z/OS databases, duplicating data in IMS DL/I and DB2 for z/OS relational databases might be required.

For example, your DB2 for z/OS application programs, written in Structured Query Language (SQL), might require data from the IMS DB database. You might be converting your site to DB2 for z/OS on a gradual basis, or you might want to take advantage of the relational technology of DB2 for z/OS for some of your IMS data.

To duplicate data between the two types of databases, you must ensure that each update to data segments occurs in both databases in a timely manner. The process of duplicating updates from an IMS DB database to a DB2 for z/OS database is known as data propagation.

The two ways to propagate data from IMS DB to DB2 for z/OS are:

- IMS DataPropagator, an IBM licensed program that provides support for data propagation and for exit routines.
- Data Capture exit routine. This is an exit routine that you write to establish a routine for data propagation. It can be written in assembler language, C language, COBOL, or PL/I, and it is called by an application program that requires data propagation.

**Note:** The Data Capture exit routine is not available to CICS. DBCTL can use the exit routine, but only for BMPs.

#### **Related reading:**

- For information about system requirements for data propagation, see *IMS Version 11 System Definition*.
- For information on the database considerations associated with the Data Capture exit routine, *IMS Version 11 Database Administration*.
- For information on writing a Data Capture exit routine, see *IMS Version 11 Exit Routines*.
- For more information about IMS DataPropagator, see:
  - *IMS DataPropagator for z/OS: An Introduction*
  - *IMS DataPropagator for z/OS Concepts*

## **HALDB Partition Selection exit routines**

Any installation-defined High Availability Large Database (HALDB) Partition Selection exit routines must be assembled and bound into a load library.

The library you choose needs to be part of the IMS SDFSRESL concatenation. Also, all HALDB Partition Selection exit routines must be linked as re-entrant.

You can specify the name of the HALDB Partition Selection exit routine during DBD generation, with the HALDB Partition Definition utility, or on the DBRC INIT.DB command.

If a HALDB Partition Selection exit routine is defined, you must also consider the Partition String Value parameter defined by the HALDB Partitioned Definition utility or the KEYSTRNG parameter on the DBRC INIT.PART command. This parameter provides the user exit routine with optional data that can be used for HALDB partition selection. The parameter is not required and may be omitted.

#### **Related reference:**

 HALDB Partition Selection exit routine (DFSPSE00) (Exit Routines)

---

## The z/OS Automatic Restart Manager (ARM)

You can use the z/OS Automatic Restart Manager (ARM) to restart a subsystem (or job) after a z/OS hardware or software failure.

In addition, in the event of a z/OS hardware or software failure that requires you to move a subsystem from one z/OS system to another, ARM will move all the subsystems defined in the same restart group as a group to a remaining z/OS system.

IMS supports ARM in these environments:

- TM-DB
- DCCTL
- DBCTL
- XRF
- FDBR

DL/I, DBB, and IMS utilities are not supported. The IMS control region is the only region restarted by ARM.

**Attention:** The DL/I SAS and DBRC regions are started internally by the IMS control region. IMS dependent regions are not automatically restarted, because they are normally restarted after the IMS control region has restarted.

The element name that IMS uses on the registration call to ARM is the IMSID. The element type is SYSIMS. Duplicate element names are not allowed by ARM. When ARM is used, the IMSIDs of online systems and FDBR systems must be unique.

ARM provides a default ARM level of 1 for SYSIMS.

The IMSID must be unique across the sysplex. ARM tries to move IMS to a surviving z/OS if a failure occurs on the z/OS or the CPC on which the IMS is executing. If the IMSID is not unique, ARM might move the IMS from the failing CPC to one that already has an IMS with the same IMSID.

If IMS is canceled by z/OS, IMS is only automatically restarted by ARM if the ARMRESTART option is specified on the CANCEL or FORCE command.

IMS maintains the following user abend table and de-registers from ARM any time one of these abends occurs:

1. U0020: USER 20 - MODIFY
2. U0028: USER 28 - /CHE ABDUMP
3. U0604: USER 604 - /SWITCH
4. U0758: USER 758 - QUEUES FULL
5. U0759: USER 759 - QUEUE I/O ERROR
6. U2476: USER 2476 - CICS TAKEOVER

The first three of these abends are the result of operator intervention. The last three abends require some external changes before IMS can be restarted.



---

## Chapter 3. Documenting the IMS system

When planning to support the administration of an online IMS system, you must consider several responsibilities that involve documentation.

Reviewing this documentation for application requirements is a necessary task when you are designing the IMS online system or responding to required changes in that design.

---

### Extracting requirements for the IMS system

Analysis of the scope and impact of an application in the online environment occurs during the design phase, as well as when application changes are proposed.

You must assess the detail of the application requirements by reviewing the following sources:

- Program specifications and logic
- Implementation plan
- Summary of business requirements
- Design change requests

When you examine application documents, you must extract several kinds of information:

- Requirements for IMS function
- Database requirements
- Predictions of the application workload
- Network definition requirements
- Security considerations
- Operating requirements
- Audit and history recommendations
- Performance factors
- Terminal requirements
- FPBUF requirements

If your system must use intelligent remote stations, you must:

- Select features of IMS that can be used to support distributed application processing (for example, ISC, MSC, LU 6.2).
- Identify the terminal support provided by IMS (including ETO).
- Identify the Fast Path requirements of your system.
- Evaluate the off-loading of application requirements to intelligent remote stations and the use of special components or screen formatting.
- Assist in the design of programs that reside in intelligent remote stations and communicate with IMS; identify the Systems Network Architecture (SNA) protocol to be used by IMS and the intelligent remote station.



---

## Participating in design reviews

As the administrator of the online IMS system, you are concerned with adequate detail in the specifications. You need to plan for, and subsequently specify, the online system. As the development of an application package progresses, reviews of the design should be held.

The following table summarizes the kind of information that you must gather and how that information relates to system administration tasks.

In addition to application design specifications, the application development team might maintain a controlling document for schedules and responsibilities. You must contribute to this plan with your own requirements and milestones, such as completion dates and testing dates for the operating procedures.

*Table 3. Administration's use of design reviews*

| Design stage                       | Information needed                                                                                                                                                                                                                                                              | Administrative tasks                                                                                                                                                                                                                                                                                                           |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Design review 1                    | <ul style="list-style-type: none"><li>• Scope of project</li><li>• Hardware and software requirements</li><li>• End-user and development contacts</li></ul>                                                                                                                     | <ul style="list-style-type: none"><li>• Analyze IMS function requirements</li><li>• Contribute to documentation plans</li><li>• Check standards compliance</li><li>• Assess network impact</li><li>• Assess DP operations impact</li><li>• Make workload predictions</li></ul>                                                 |
| Design review 2                    | <ul style="list-style-type: none"><li>• Use of MFS and screen usage characteristics</li><li>• Elements of end-user control</li><li>• Network planning</li><li>• Pointers for operator control</li><li>• Transaction workload</li></ul>                                          | <ul style="list-style-type: none"><li>• Establish MFS library control, identify format names</li><li>• Coordinate dictionary use</li><li>• Check naming standards</li><li>• Track network requirements</li><li>• Plan security strategy</li><li>• Begin RTO and MTO procedures</li><li>• Predict processing workload</li></ul> |
| Design review 3                    | <ul style="list-style-type: none"><li>• Message definition</li><li>• Conversational attributes</li><li>• Databases and programs</li><li>• System resource requirements</li><li>• Need for message edit</li><li>• Recovery considerations</li><li>• Security and audit</li></ul> | <ul style="list-style-type: none"><li>• Calculate message queues</li><li>• Calculate SPA data</li><li>• Specify system and JCL</li><li>• Finalize system requirements and hardware plan</li><li>• Specify message edit coding</li><li>• Begin recovery procedures</li><li>• Develop security design</li></ul>                  |
| Database application design review | <ul style="list-style-type: none"><li>• Database maintenance</li><li>• Database sharing</li><li>• Validation and acceptance plans</li><li>• Performance predictions</li><li>• Monitor plans</li></ul>                                                                           | <ul style="list-style-type: none"><li>• Document online databases and image copy requirements</li><li>• Choose system integrity options</li><li>• Establish system availability</li><li>• Establish performance criteria and plan monitoring</li></ul>                                                                         |
| Logic review                       | <ul style="list-style-type: none"><li>• Monitor pointers</li><li>• Program preload</li><li>• Virtual storage needs</li><li>• Database processing intent conflicts</li></ul>                                                                                                     | <ul style="list-style-type: none"><li>• Develop a monitoring strategy</li><li>• Plan dependent regions</li><li>• Develop scheduling algorithm</li><li>• Estimate buffer pool and system data set resources</li></ul>                                                                                                           |

**Related reading:** For more information on the purpose and scope of design reviews, see *IMS Version 11 Database Administration*.

---

## Establishing naming conventions

A critical part of the application specification and the control of the IMS online system design is maintaining naming conventions for your resources.

When you define a large system that has many resources, the ability to recognize the characteristics of the resource by its name has many advantages:

- The system definition input is easier to check, and the identification of changes is easier.
- The MTO control is more effective and efficient and less prone to error.
- The modification of the application design can more easily recognize already defined resources rather than creating ambiguous or unnecessary additional resources.

You should establish naming conventions, in cooperation with database administrators, for at least the following resources:

- Databases, their DD names and data set names
- Image copy and change accumulation data set names
- Segment and field names
- PSB and program names
- Transaction codes
- MFS format names
- LTERM and node names
- ETO terminal and user names
- LU 6.2 descriptor names
- Online log data set names
- System log data set names
- IMS Monitor output data set names
- Link names and IMS system IDs (for Multiple Systems Coupling)
- Fast DB Recovery region names
- High Availability Large Databases (HALDBs), and HALDB partitions, DD names, and data set names.

The following table shows some examples of naming conventions that can be applied to resources controlled by IMS for online applications.

*Table 4. Examples of naming conventions*

| Resource    | Naming Convention | Description |                                             |
|-------------|-------------------|-------------|---------------------------------------------|
| Transaction | Taaatsss          | <b>T</b>    | Transaction                                 |
|             |                   | <b>aaa</b>  | Application identifier                      |
|             |                   | <b>t</b>    | U for update, or R for inquiry transactions |
|             |                   | <b>sss</b>  | Transaction sequence                        |

Table 4. Examples of naming conventions (continued)

| Resource                           | Naming Convention | Description |                                                                           |
|------------------------------------|-------------------|-------------|---------------------------------------------------------------------------|
| LTERM name                         | cnnxiiii          | <b>c</b>    | L for local, S for switched, N for non-switched                           |
|                                    |                   | <b>nn</b>   | A 2-character code for terminal type                                      |
|                                    |                   | <b>x</b>    | A 1-character attribute indicating the screen size, printer, or component |
|                                    |                   | <b>iiii</b> | A 4-character identifier                                                  |
| MFS<br>(MSG name)<br>(MID and MOD) | aaaiiii           | <b>aaa</b>  | Application identifier                                                    |
|                                    |                   | <b>iiii</b> | A 4-character identifier                                                  |
| MFS<br>(FMT name)<br>(DIF and DOF) | aaaiii            | <b>aaa</b>  | Application identifier                                                    |
|                                    |                   | <b>iii</b>  | A 3-character format identifier                                           |
| Module name                        | Maaaiiii          | <b>M</b>    | Module name                                                               |
|                                    |                   | <b>aaa</b>  | Application identifier                                                    |
|                                    |                   | <b>iiii</b> | A 4-character identifier                                                  |
| Job name                           | Jaaannnn          | <b>J</b>    | Job name                                                                  |
|                                    |                   | <b>aaa</b>  | Application identifier                                                    |
|                                    |                   | <b>nnnn</b> | A 4-character job identifier                                              |

For more information on naming conventions, see the:

- *IMS Version 11 Database Administration* for recommendations for naming conventions for databases, PSBs, programs, and HALDBs.
- *IMS Version 11 System Definition* for a list of restricted names.
- *IMS Version 11 Communications and Connections* for specific naming conventions for ETO.

## Using a data dictionary

If you use the IBM DB/DC Data Dictionary licensed program (program number 5740-XXF), you can use several of its features to assist in system and network documentation. The product supports interactive processing, enabling you to use a terminal to create the documentation at your convenience.

The product provides standard categories for databases, segments, fields, PSBs, PCBs, transactions, programs, and the system. You can use these categories to build up detailed descriptions of the online IMS system resources. Procedures and execution JCL can also be recorded (as User Data). Three of the stage 1 macros—APPLCTN, DATABASE, and TRANSACT—can be produced as card output.

You can use the Extensibility Feature to define terminal subjects and specify a selection of attributes to record for each terminal. With the Description and User Data segments, other free-form text can be associated with the terminal subjects.

## Documenting the system characteristics

You must create an independent body of information that documents the design and operation of the production system. This material is derived from the application requirements and includes the intended use of IMS facilities.

The following list describes the main documentation activities.

- Documenting the IMS system definition

The detail of your design for an IMS online system is reflected in the system definition macro specifications. The input to the first stage of system definition processing can be used as a major documentation tool. Including comments on the choice of parameters should help you control definition changes.

The IMSCTRL macro statement offers an ETO option that creates a report of current ETO descriptors in the network. You can use the ETO descriptor report to monitor your current network definition.

**Related reading:** For more information about commenting on parameters and information about the IMSCTRL macro, see *IMS Version 11 System Definition*.

- Documenting the IMS network in DB/DC and DCCTL environments

Within the broader scope of system documentation, you need to document the details of the IMS network. Start this documentation as soon as initial plans for an application system become available. The advantage to starting early is that you can become familiar with the physical network and the way it will logically be used by IMS connections. When your online IMS system uses terminals that are part of a network defined to VTAM, you might be able to use some of the documentation developed by the system programming staff.

Documenting in detail allows you to:

- Familiarize yourself with the features and operation of the terminal devices
- Find out if the application is going to use the terminal in an unusual way and, if so, research the potential problems
- Prepare for IMS system definition stage 1 input
- Prepare ETO descriptors
- Prepare LU 6.2 descriptors
- Plan for the installation and network generations
- Understand the operational aspects of subsets of the total configuration

- Documenting terminal profiles in DB/DC and DCCTL environments

One way to document the required network, after the major design is stable, is to record the intended use and characteristics of each terminal. You construct a profile that contains:

- The terminal type, its required features, and the type of connection it is to use
- What options were chosen for the device and the reason for the choice
- The characteristics of how the terminal is to be used by the application
- The proposed LTERM names that can be associated with the terminal or the user
- If a VTAM-supported device, the node name and transmission characteristics
- The extent of the proposed usage and whether the device is to be shared with non-IMS users
- The diagnostic procedures that are appropriate for the device

- User profiles and user profile security

For dynamic terminals, you should maintain a record of the characteristics of dynamic terminals and users to make future changes easier.

- Documenting transaction profile names for APPC/IMS

Definitions of transaction program names (TPNs) are contained in the APPC/MVS resource, TP\_Profile. Within TP\_Profile, TP profile data sets provide attribute information for each TPN. You can define TPNs that have different characteristics for each LU name with which they are associated.

**Related reading:** For more information on TP names, see *IMS Version 11 Communications and Connections*.

- Documenting the configuration of the production system

You should create a configuration map showing how individual terminals and clusters of terminals are to be connected in the network. This gives you some insight as to how to specify the network control to the MTO. Having the configuration map is also an aid when you interact with various technical specialists. The map should show:

- Processors or host computers
- Channels and lines (including the type of line)
- Communication controllers
- Control units and the terminal attachments
- VTAM node names
- For XRF systems, USERVAR or MNPS ACB name
- Alternative connections or configurations

To solve problems that might arise, you must be able to identify the terminal or control unit that has the problem. You can assist both end users and service personnel by placing identifying labels on the devices. In addition to the IMS address or node name and the LTERMs appropriate for that terminal, include the hardware address, circuit identification, and other data that might be necessary.

---

## Chapter 4. Introduction to IMS operations and recovery

For any large system, operations and recovery go together: *operations* refers to the day-to-day activities involved in keeping the system running smoothly, and *recovery* refers to the activities required to bring a failed system back online.

Information about operations would be incomplete if it did not describe what you need to do for recovery, when the system (or some part of it) runs into problems.

Another important task related to operations is planning for both day-to-day operations and for recovery. Part of the task of planning is developing procedures that specify the roles of the IMS operators and system programmers. These procedures clarify the responsibilities of each of these roles.

This topic introduces both operations and recovery for IMS and describes some of the tools provided by IMS and z/OS for both tasks.

---

### Understanding the operations task

The task of planning for operations has two major parts: selecting functions and tools and developing operating and end-user procedures.

First, you must decide how to use the tools IMS provides for operating your system. This includes choosing, for example, whether to use dual logging when setting up your log, how often to make backup copies of your database, and whether to use DBRC to control recovery of databases.

Second, you must develop procedures for operating and using IMS. Operating procedures must tell operators how to:

- Start and restart IMS
- Control IMS
- Make online changes to modify IMS
- Shut down IMS
- Run various IMS utilities
- Recover from IMS and other failures

These procedures are primarily for the master terminal operator (MTO), who operates IMS from the master terminal. They are secondarily for the people who assist in keeping IMS running smoothly, such as the recovery specialist who works with many of the more complex IMS recovery problems.

End-user procedures tell end users how to:

- Operate their terminal
- Establish a connection to IMS
- Communicate with IMS (using their specific applications)
- Terminate a connection to IMS
- Respond to any error conditions they encounter

This information is not an exhaustive treatment of operations. More detailed information on various aspects of operating your system is provided in other IMS topics.

**Related reading:**

- *IMS Version 11 Operations and Automation* gives guidance level information and details about the commands that are used to operate IMS.
- *IMS Version 11 Database Utilities* and *IMS Version 11 System Utilities* give guidance level information and detail about the various utilities that are available with IMS.

## **Automated operations**

Automated operations are tools and techniques that help you improve your installation's productivity. As your system grows more complex and your message traffic increases, automating certain tasks can increase your system's efficiency.

You can use automated operations to:

- Minimize errors
- Increase availability
- Expedite problem diagnosis and prevention

IMS provides the following tools to help you automate your operations:

- Time-Controlled Operations (TCO)
- Automated Operator Interface (AOI)
- REXX SPOC API

IMS operations can also be automated using Tivoli NetView for z/OS. Because Tivoli NetView for z/OS functions independently of IMS, it can gather information and issue commands that are not available to IMS.

There are many advantages for automating tasks. Many of the jobs that operators perform are simple, repetitive tasks, such as monitoring the system and issuing recovery commands. You can often automate these jobs and thereby free the operator for more complex activities, such as implementing operational procedures.

In the IMS environment, TCO and AOI can improve your operator productivity by:

- Improving operator productivity and accuracy. Automating operator tasks simplifies procedures, reduces operator input, and minimizes operator errors. For example, TCO can reduce operator input by automatically monitoring system status, starting message regions and telecommunication lines, and notifying users of system status.
- Expediting problem determination. An operator's primary job is to bypass or fix problems quickly. If adequate information about a problem is not available, it might not be possible for the operator to fix it quickly. Automated operations are especially suited for problem determination. For example, TCO can automatically gather necessary information and do diagnostic analysis so that a problem can be quickly identified and corrected.

Identifying operations that can be automated and implementing them requires a sound understanding of your installation's operations. You must collect and analyze various resources to identify which tasks are good candidates for automation. Resources you should analyze include:

- System logs

- Problem management reports
- Record of calls to the help desk
- Operators' notes

Repetitive and predictable tasks are good candidates for automation.

The process of automating operations should be an iterative one. After you have developed and used automated procedures, you should evaluate them and, in the process, consider whether any new tasks can be automated.

#### **Related reading:**

- For more information about automation tools provided with IMS, see *IMS Version 11 Operations and Automation*.
- For more information about Tivoli NetView for z/OS, see *Tivoli NetView for z/OS Installation: Getting Started* or the *Tivoli NetView for z/OS User's Guide*.

---

## **Understanding the recovery task**

Recovery is the task of restoring a failed system, application program, or database back to normal operations.

Recovery is the largest and most complex part of operations. For this reason, any discussion of operations is primarily a discussion of recovery.

Recovering a system can involve:

- Databases
- User requests to process data
- Programs that do the processing (application programs)
- Output sent to users

A recoverable system ensures:

- Data is not lost
- Incomplete changes to a database are not saved

Data can be lost in two ways. It can be physically lost—the disk or tape on which it resides can be damaged or misplaced. Or it can be “logically” lost—the data becomes incorrect or loses its relationship to other data. A system that ensures *data integrity* ensures that data cannot be lost or saved incompletely.

### **Example: a system without recovery**

To show how things can go wrong, consider a hypothetical system that has no recovery mechanisms. A user requests a transaction that handles a customer order.

Part of the application program's work is to:

- Decrease the stock-on-hand number in the database
- Add the cost to the customer's outstanding bill
- Add the amount to a daily-total-sales accumulator

The following list provides different failure scenarios:

- Effects of system failure

What if the system fails between the time the stock-on-hand number is decreased and the charges are allocated? If the user thinks the transaction is



complete, the customer bill record and the total sales record in the database will be incorrect. If the user thinks the transaction failed and reruns it, the stock-on-hand number will be decreased a second time, and that record will also be wrong. In either case, data integrity has been lost.

- Effects of programabend

What if the program begins changing records in the database and then terminates abnormally (abends)? The result is the same as with the effects of system failure: a half-changed database. A half-changed database is one whose integrity has been lost.

- Effects of an I/O error

Suppose the program tries to read the stock-on-hand record and encounters a device error. Because of some I/O problem, the record cannot be obtained. Therefore, the program cannot run, and the customer order cannot be filled. In addition, any other work that depends on this record can no longer be completed.

Also, what if the program has already changed some records expecting that it could read and update this other record? Again, the overall integrity of the data has been lost, because it is only partially updated.

- Effects of queue loss

In large systems, work requests are often saved (queued) and processed later, when the workload permits. Output that the programs send back to the requesters (or others) is also often queued and sent later, when the workload permits.

If the system fails between the time a request is entered (input) and the time it is taken off the queue and executed, the request might be lost. The same is true for queued output: If the system fails between the time the application program supplies its output and the time that output can be taken off the queue and actually sent, those results might be lost.

## Requirements for a recoverable system

In order for a system to be recoverable, it must protect the database from half-complete (and, therefore, incorrect) changes. It must be able to deal with the physical loss—unavailability or disappearance, in part or in whole—of the database. Also, it must protect the input and output queues from being lost.

## Steps in the process of recovery

The process of recovery includes two major tasks: planning what you intend to do for recovery before a problem occurs and reacting to a problem after it has occurred to fix the problem.

The task of planning for IMS recovery can be further refined:

- Set up logging
- Establish checkpoints
- Make backup copies
- Create procedures and assign responsibilities

The task of reacting can also be further refined:

- Notice error
- Shut down IMS (only if necessary) — can be a full or partial shutdown
- Diagnose and fix (or possibly bypass) error
- Use recovery utilities and services (if necessary) — can be run online or offline

- Restart IMS (or the now fixed part of IMS)

Of course, real recovery is rarely as straightforward as these two lists suggest.

## The mechanisms of recovery in IMS

IMS is designed to be recoverable and has a number of mechanisms to help you in these areas. Some function automatically; others require your involvement.

Successful recovery depends on two things:

- Having a safe point to return to, a known point of integrity. This is a point at which you know data is correct and make a record of it.
- Keeping track of what processing has been done since that safe point.

If these two types of information are gathered, recovery is almost always possible.

### Synchronization points

Strictly speaking, all you need to be able to restart a system after a failure is the log: as long as you start out correctly and log everything you do thereafter, you can always recover. Restart in this case would require the system to read every log record ever written for the system.

The more time that has passed since the first time you started the system, the more log records will have been written, and the more time-consuming the restart will be.

**Definition:** A synchronization point, or *sync point*, is a known point of data or system integrity. It is a point in time at which all of the changes made to the database are complete. If a failure makes restart necessary, you can begin reprocessing from the most recent sync point.

The value of having a sync point is that when IMS restarts, it can ignore all logs created before the sync point was taken. If your sync points are infrequent, they become less valuable because the amount of log data IMS has to read during a restart increases. Thus, you should establish periodic sync points. Then, if a problem arises, you do not need to examine as much of the log, because you established the last sync point relatively recently.

IMS uses two types of sync points: those taken by IMS itself (called *system checkpoints*), and those taken by individual application programs running under IMS (called application program sync points or application program *commit points*).

### System checkpoints:

IMS automatically takes periodic system checkpoints, so if it is necessary to restart IMS, it can begin at the last checkpoint (or an earlier one if the work done before the checkpoint was not complete).

The interval between system checkpoints is measured by how many log records IMS writes, so the checkpoints are a measure of activity rather than elapsed time. You select the interval between system checkpoints when you install IMS.

As with most choices involving preparation for recovery, the impact of taking frequent checkpoints must be weighed against the advantages of faster recovery. Because IMS does not suspend work to take a checkpoint, decreased transaction speed is not a significant impact factor.

**Recommendation:** Do not suppress system checkpoints to improve system performance because emergency restart must always go back at least two system checkpoints (or back to a prior restart).

The IMS MTO can also request an additional system checkpoint at any time by issuing a `/CHECKPOINT` command.

IMS takes a system checkpoint when any of the following occur:

- IMS starts
- A regular interval has elapsed
- The MTO enters a command (`/DBRECOVERY`, `UPDATE DB STOP(ACCESS) OPTION(FEOV)`, `/CHECKPOINT`, or `/SWITCH OLDS CHECKPOINT`)
- IMS shuts down

### Application sync points:

IMS system checkpoints help you recover the IMS subsystem, but you also need to be able to restart application programs if they fail. IMS allows application programs, both batch and online, to take sync points.

An application program sync point has two purposes:

- It marks an intermediate completion point—a place at which the work finished so far is judged to be correct. Any future recovery can take place from this point.
- It frees locks held on database records the program has updated and enqueues any output messages the program created. IMS can send enqueued messages to their destinations.

Application program sync points are sometimes called *commit points*. By taking a sync point, the program is committing itself to the accuracy and completeness of what it has done, and releasing the data for use by other applications.

## Logging

Logs are lists of activities, lists of work that have been done, and lists of changes that have been made. By knowing the state of the system when things were all right, and knowing what the system has done since then, you can recover it in the event of a failure.

Logs make recovery and restart possible. As IMS operates, it constantly writes its event information in log records. The log records contain information about:

- When IMS starts up and shuts down
- When programs start and terminate
- Changes made to the database
- Transaction requests received and responses sent
- Application program checkpoints
- System checkpoints

The following figure illustrates how logging works in an IMS online environment; logging in a batch environment is simpler.




Figure 22. Overview of the logging process

IMS externalizes log records by writing them to an Online Log Data Set (OLDS). To enhance performance and to optimize space on OLDSs, incomplete or partially filled buffers are written to a Write Ahead Data Set (WADS) when necessary for recoverability. The WADSs are high-speed DASD data sets with a high write rate. Only complete log buffers are written to OLDSs. When the log data is on an OLDS, the equivalent WADS records are ignored.

IMS uses a set of OLDSs in a cyclical way, which allows IMS to continue logging when an individual OLDS is filled. Also, if an I/O error occurs while writing to an OLDS, IMS can continue logging by isolating the defective OLDS and switching to another one. Once an OLDS has been used, it is available for archiving to a System Log Data Set (SLDS) on DASD or tape by the IMS Log Archive utility. The utility can be executed automatically through an IMS startup parameter (ARC=).

When IMS is close to filling the last available OLDS, it warns you so you can ensure that archiving completes for used OLDSs or add new OLDSs to the system. You can also manually archive the OLDSs to SLDSs using the IMS Log Archive utility. An SLDS can reside on DASD or tape. After an OLDS is archived, it can be reused for new log data. You use SLDSs as input to the database recovery process.

When you archive an OLDS, you can request that IMS write a subset of the log records from the OLDS to another log data set called the recovery log data set (RLDS). An RLDS contains only those log records required for database recovery.

While IMS is running and logging its activities, it takes periodic system checkpoints, and writes the checkpoint notification to another data set called the restart data set (RDS). IMS uses the RDS when it restarts to determine the correct checkpoint from which to restart.

In a batch environment, IMS writes log records directly to an SLDS, and does not use either the OLDSs or WADSs. Just as in the online environment, batch SLDSs can reside on DASD, tape, or mass storage. If the SLDS is on DASD, you can use an IMS utility to copy log records from DASD either to tape or to other DASD.

### The IMS log data sets:

These topics describe the OLDS, WADS, SLDS, RLDS, restart data set (RDS), RECON data set, the z/OS log data set used for CQS, the CQS system checkpoint data set, and the CQS structure checkpoint data set.

#### *Online log data set:*

IMS uses the OLDS only in the online environment. The OLDS contains all the log records required for restart, recovery, and both batch and dynamic backout. The OLDS holds the log records until IMS archives them to the SLDS.

Define all of the OLDSs in the IMS procedure library (IMS.PROCLIB) using the OLDSDEF statement. The OLDS must be preallocated on a direct-access device. You can also dynamically allocate additional OLDSs while IMS is running by using the /START OLDS command.

IMS uses the Basic Sequential Access Method (BSAM) to write log records to the OLDS, and the Overflow Sequential Access Method (OSAM) to read the OLDS when IMS performs dynamic backout. Although referred to as a data set, the OLDS is actually made up of multiple data sets that wrap around, one to the other. You must allocate at least three, but no more than 100, data sets for the OLDS.

**Related reading:** For more information about allocating data sets for the OLDS, see “Allocating data sets” in *IMS Version 11 System Definition*.

You can specify that the OLDS use *dual logging*, which is the duplication of information on two logs. When you use dual logging, an I/O error on either the primary or secondary data set causes IMS to close the nonerror OLDS and mark the error OLDS in the Recovery Control (RECON) data set as having an I/O error and a close error. IMS then continues logging with the next available pair. For dual logging, the minimum number of data sets is three pairs, and the maximum number is 100 pairs.

IMS uses as many OLDSs as you allocate. IMS issues a message each time it changes the current OLDS. This message identifies the OLDS being closed and the next OLDS to be used.

When any of the following events occur:

- IMS fills an OLDS
- An I/O error occurs
- You issue one of the following commands:
  - /DBDUMP DB
  - /DBRECOVERY DB
  - UPDATE DB STOP(UPDATES) OPTION(FEOV)
  - UPDATE DB STOP(ACCESS) OPTION(FEOV)

IMS does the following:

- Opens the next OLDS

- Notifies DBRC and the MTO that it is using a new OLDS
- Closes the current OLDS (both primary and secondary if you are using dual logging)

When IMS is using the last available OLDS, it alerts the MTO that no additional OLDS space is available. If archiving has not finished by the time all of the OLDSs are full, IMS waits until OLDS space becomes available. IMS will not log to an OLDS containing active data that is not yet archived. You must run the Log Archive utility to free the OLDS space. After IMS uses the last allocated OLDS, it reuses the first OLDS, if it has been archived.

You can use the `/STOP` command to stop and dynamically deallocate an OLDS. When stopped, that OLDS is no longer involved in the wraparound process.

**Recommendation:** Stop an OLDS when an error occurs that requires that OLDS to be recovered.

**Restriction:** You cannot stop the current OLDS. You cannot stop any OLDS when two or fewer OLDSs are currently available.

Similarly, you can use the `/START` command to start and dynamically allocate an OLDS. IMS retains the status of an OLDS (in-use, stopped, and so on) from one restart to the next.

Jobs to archive OLDSs might not complete in the order in which the OLDSs were created. For example, one OLDS might not yet be archived, but a subsequent OLDS might already be archived. When this occurs, IMS issues message DFS3259I and uses the next available OLDS.

The DBRC RECON data set contains information about the OLDSs for each IMS subsystem. Information in the RECON data set indicates whether an OLDS is available for use or contains active log data that must be archived.

#### *Write-ahead data set:*

IMS uses the write-ahead data set (WADS) only in the online environment. The WADS contains a copy of committed log records that are in OLDS buffers, but that have not yet been written to the OLDS.

In order to maximize log efficiency, IMS uses a log write-ahead function to write partially filled, padded blocks to the WADS (rather than the OLDS). IMS continually reuses WADS space after writing the appropriate log data to the OLDS.

The log write-ahead function ensures that all log records are on the log before IMS writes changes to a database. IMS updates a database in any of the following situations:

- When IMS needs to reuse the database buffer (if this is before commit)
- During commit
- During VSAM background write

If IMS fails, you use the log data in the WADS to complete the content of the OLDS and then close the OLDS as part of an IMS emergency restart or as an option of the Log Recovery utility. If you close the OLDS during emergency restart, you must include the WADS in use at the time of the failure.

You must preallocate and format the WADS on a DASD device that supports count key data (CKD) architecture. Format a WADS using the `FORMAT WADS | ALL` keywords on either the `/NRESTART` or `/ERESTART` commands. All WADSs must be on the same device type and should have the same space allocation. You can also dynamically allocate additional WADSs using the `/START WADS` command.

You can change any of the following specifications for the WADS during an IMS restart:

- Number of WADSs
- Sequence of WADSs
- WADS names
- Use of single or dual WADSs

**Recommendation:** To eliminate potential resource contention, place the WADS on a low-use device that is different from the device you use for the OLDS.

If you place the WADS on the same device as one of your OLDSs and use full-track blocking for the OLDS (in which a block is equal to a full track), the device should be able to handle infrequent OLDS seeks. Contention can still occur.

If the WADS and OLDS are on the same device, the Log Archive utility (DFSUARC0) or dynamic backout can cause severe contention between an OLDS being archived and an active WADS.

*System log data set:*

IMS uses the SLDS in both the online and batch environments. In the online environment, an SLDS contains archived OLDS data. In the batch environment, an SLDS contains current log data.

Each execution of the Log Archive utility creates an SLDS. One SLDS can contain data from one or more OLDSs. You use an SLDS as input to the database recovery utilities (Database Recovery, Database Change Accumulation, and Batch Backout). You can also use an SLDS during an emergency restart of IMS. SLDSs can be stored on DASD, tape, or other mass storage.

DBRC maintains information about SLDSs in the RECON data set:

- For batch subsystems, DBRC maintains SLDS information in the PRILOG and SECLOG records.
- For online subsystems, DBRC maintains SLDS information in the PRILOG and SECLOG records only if you do not specify an RLDS when you run the Log Archive utility. Otherwise, DBRC maintains SLDS information in PRISLD and SECSLD records.

The Log Archive utility tells DBRC which OLDS it is archiving and which SLDS it is creating. The IMS online system can reuse OLDSs that have been archived.

Generally, you want to copy all the log records from the OLDS to the SLDS, but you can specify specific records. If you want to omit some types of log records from the SLDS in order to save space, include the `NOLOG` keyword when you run the Log Archive utility. The SLDS must always contain those records that might be needed for database recovery, batch backout, or IMS restart. The records that you can omit are:

**X'10'** Security violation records



- X'45' Statistics records written during checkpoint
- X'5F' Call trace record
- X'67' Communications (SNAP) trace records
- X'69' Unauthorized ID record (for 3275 display terminal)

IMS dynamically allocates an SLDS during IMS restart whenever log data required for restart read processing is not available from an OLDS. The OLDS might be unavailable because it has been archived, and because one of the following is true:

- The OLDS has been reused.
- The PRIOLDS and SECOLDS records have been deleted from the RECON data set.

To allow IMS to dynamically allocate SLDSs, you must specify the SLDS device type Dynamic Allocation macro (DFSMDA). DBRC provides the data set name and volume information required for dynamic allocation.

#### *Recovery log data set:*

When you run the Log Archive utility to create an SLDS, you can also request creation of an recovery log data set (RLDS). The RLDS can be stored on DASD, tape, or other mass storage device.

The RLDS contains only the log records needed for database recovery:

- X'24' Database error records
- X'3730' Sync point records
- X'4001' Checkpoint records
- X'4084' Checkpoint records that contain Fast Path DMCB/DMAC control blocks.
- X'4098' Checkpoint records for the end of Fast Path checkpoints.
- X'42' Checkpoint ID records
- X'5612' End of phase 2 commit records
- X'5701' Database begin update records
- X'59' Fast Path database change records
- X'505x' Database change records

IMS maintains RLDS information in the RECON data set in the PRILOG and SECLOG records. Whenever possible, DBRC uses the RLDS in place of SLDSs when creating JCL for the Database Recovery and Database Change Accumulation utilities. Using the RLDS rather than the SLDS is more efficient because the RLDS contains less information than the SLDS.



#### *Restart data set:*

IMS writes system checkpoint information to the restart data set (RDS). During each checkpoint, IMS creates or updates a checkpoint ID table; IMS uses this table during IMS restart to determine from which checkpoint to restart the system.

If, for any reason, the RDS is not available at restart, IMS can obtain the required checkpoint information from the log. However, using only the log could considerably lengthen the restart process.

Generally, you do not need to know the content of the RDS. IMS finds the information it needs in the RDS and uses it automatically during a restart.

#### *RECON data sets:*

DBRC automatically records information in the RECON data sets. Because both RECON data sets contain identical information, this information refers to them as a single data set.

IMS uses the RECON data set in many situations:

- During warm start and normal and emergency restarts. The RECON data set shows which data set—OLDS or SLDS—contains the most recent log data for each DBDS that you registered with DBRC.
- During logging, the RECON shows its latest status for the OLDS and whether the OLDS has been archived.
- For a recovery utility, DBRC selects the correct data sets.

#### *z/OS log data set:*

The IMS Common Queue Server (CQS) records information about the data in the IMS shared queues in the z/OS log data set. The z/OS system logger serves the same purpose for CQS as the OLDS serves for IMS: it records all necessary information so CQS can recover structures in the coupling facility and restart after failure.

CQS writes log records for each pair of coupling facility list structures to a separate log stream. This log stream is shared among all the CQS subsystems that share the structure pair. z/OS merges the log streams to make recovery possible.

**Related reading:** For more information on the z/OS system logger, see “Using System Logger Services” in *z/OS MVS Programming: Assembler Services Guide*.

#### *CQS system checkpoint data set:*

Each CQS subsystem maintains a system checkpoint data set for each structure pair in the coupling facility.

Whenever the CQS subsystem takes a system checkpoint, it writes some control information to this data set. It also writes log records to the system logger log stream.

The system checkpoint data sets are not shared among CQS subsystems.

#### **Related concepts:**

“Using CQS system checkpoint” on page 180

### *CQS structure recovery data set:*

Whenever a CQS subsystem takes a structure checkpoint, it writes a snapshot of the message queues to a structure recovery data set. It also writes some log records to the system logger log stream. The structure recovery data set is used to recover the message queues.

CQS subsystems in a sysplex share the structure recovery data sets; there is one pair (two data sets) for each structure. CQS alternates between the two for each checkpoint.

### **Related concepts:**

“CQS structure recovery” on page 187

### **Archiving log records:**

For online systems, you can automatically or manually initiate archiving log records from the OLDS to the SLDS.

### *Automatic archiving:*

If you have a large system with a lot of activity, you can minimize your intervention in the archiving process by using automatic archiving. It will eliminate the need to continually monitor logging to determine when to archive.

By default, IMS archives each OLDS when it is full. However, you can control how often archiving occurs by specifying how many OLDSs must be full before IMS archives them. Use the ARC= execution parameter or the AUTOARCH keyword of the /START command to control automatic archiving.

**Recommendation:** You must archive an OLDS before IMS can reuse it; be sure to archive frequently enough to avoid running out of OLDS space. If you run out of OLDS space, IMS waits until OLDS space becomes available.

**Related reading:** For more information about the ARC= execution parameter, see *IMS Version 11 System Definition*. For more information about the AUTOARCH keyword of the /START command, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

### *Manual archiving:*

If you archive infrequently or at irregular intervals, you can initiate archiving yourself. Use the Log Archive utility (DFSUARC0) to archive logs manually.

You can use the DBRC GENJCL.ARCHIVE command to produce JCL for the Log Archive utility. You can issue this command using the DBRC Command request, the Database Recovery Control utility (DSPURX00), or an IMS online command (/RMGENJCL).

The SLDS can be on DASD or tape. Archiving is also useful for batch systems to free disk space if your SLDSs are on disk. Use the Log Archive utility to copy an SLDS from DASD to tape. Because DASD and tape typically have different block sizes, the utility reblocks the log records while it copies them.

IMS notifies DBRC whenever an OLDS is either filled or closed or both. DBRC updates the RECON data set to indicate that the OLDS is available to be archived.

Using the Log Archive utility, you can archive multiple OLDSDs to a single SLDS as long as the OLDSDs being archived were created consecutively by IMS. The JCL supplied to the utility defines which and how many OLDSDs are to be archived. The GENJCL facility of DBRC allows you to specify:

- Which OLDSDs should be included in the created JCL
- That all OLDSDs not yet archived should be included

If all the specified OLDSDs are archived successfully, DBRC updates the RECON data set to indicate that the OLDSDs are now available for reuse by the online system. If the Log Archive utility job fails, re-run it.

If you do not specify automatic archiving, you must create the JCL to run the utility. If you specified automatic archiving, IMS calls the DBRC GENJCL function to generate JCL for the utility when the specified number of OLDSDs has been filled or closed.

If the DBRC JCLOUT DD statement for the GENJCL output is directed to the internal reader, the archive jobs are automatically started. The following figure shows an overview of the Log Archive utility.




Figure 23. Overview of the Log Archive utility

#### Copying an SLDS or an RLDS:

You can use the Log Archive utility to copy an SLDS or RLDS to a new data set; however, you cannot use the GENJCL.ARCHIVE command to generate JCL to copy these data sets.

You can also use the Log Archive utility to create an RLDS or user data set. Use the Log Archive COPY control statement to do any of the following:

- Specify the user data sets to which you want log records copied.
- Determine which log records should be copied to a user data set.
- Specify that all log records required for database recovery should be copied to an RLDS.

#### *Customizing archiving:*

You can write user exit routines to process log records and copy certain log records to user data sets. For example, you can copy all records required for restarting Batch Message Processing programs (BMPs) to a user data set.

To customize archiving, specify the entry points for the exit routine using Log Archive utility control statements. IMS gives control to the exit routines when:

- The Log Archive utility is initialized.
- IMS reads the OLDS.
- The Log Archive utility terminates.

#### **Related concepts:**

 Guidelines for writing IMS exit routines (Exit Routines)

#### **Related reference:**

 Log Archive utility (DFSUARC0) (System Utilities)

#### **Tracing the log:**

You can trace logging activity to diagnose performance problems or problems with IMS. As with all traces, the trade-off when using log tracing is between increased diagnostic capabilities and the overhead of running the trace.

The overhead of tracing the log can be greater than running other types of trace, especially if you request an external trace, that is, if the trace itself is logged. But you can write the external trace to an external trace data set, and add no extra burden to the OLDS.

Specify log tracing in one of the following ways:

- Use the DLOG parameter on the OPTIONS control statement when you initialize IMS.
- Use the /TRACE command and omit the DLOG keyword. You can turn log tracing on and off and control whether it is to be logged to the OLDS or to an external trace data set.

#### **Reducing Fast Path logging:**

Because IMS holds updates for DEDBs in storage before writing entire VSAM control intervals (CIs), you can reduce the logging for Fast Path data. You reduce the log volume by logging only the changed data for each log record during replace (REPL) calls. You can reduce the logging only if the length of the segment remains unchanged.

Use the LGNR parameter of the IMS or DBC procedures to determine the maximum number of Fast Path DEDB buffer alterations that are to be held before IMS logs the entire VSAM control interval (CI). Use the Fast Path Log Analysis utility (DBFULTA0) to evaluate the value you should use for the LGNR parameter.

**Related concepts:**

-  Fast Path EXEC parameters in DBCTL (System Definition)
-  Fast Path EXEC parameters in DCCTL or DB/DC (System Definition)

**Related reference:**

-  Fast Path Log Analysis utility (DBFULTA0) (System Utilities)

**Using DBRC to track batch job logs:**

An IMS online subsystem always uses DBRC for tracking logs, but a batch job need not use DBRC. If you use DBRC for batch jobs, DBRC tracks which batch jobs create which SLDS.

**Recommendation:** Use DBRC for batch jobs to eliminate the need to manually keep track of batch SLDSs.

You do not need to create a log for read only (PROCOPT=G) batch jobs, but you do need to create a log for update jobs. For update jobs using DBRC, you cannot use DD NULLFILE or DD DUMMY in the JCL for the log data set.

Specify the use of DBRC during IMS system definition by using the DBRC keyword in the IMSCTRL macro. While IMS is running, you can use the DBRC= execution parameter in the DBBBATCH and DLIBATCH procedures to override the value specified during system definition. If you specify the FORCE keyword during system definition, you must use DBRC, except when you run the Log Archive (batch only), Log Recovery, or Batch Backout utilities.

**Consolidating logged database changes for recovery:**

You can use the Database Change Accumulation utility (DFSUCUM0) to extract and consolidate database change records from SLDSs or RLDSs. The Database Change Accumulation utility optimizes the database change records for recovery and saves them to a change accumulation data set. You can use the change accumulation data set as input to the Database Recovery utility (DFSURDB0).

As IMS runs, the number of SLDSs or RLDSs increases. You can use these data sets to recover a lost or damaged database, but to use them without change would be inefficient for the following reasons:

- Each SLDS or RLDS contains a record of activities of the entire IMS subsystem and of all the data sets for all the databases. Yet when you are recovering a database, you usually only recover a single data set. Thus, much of what is in the SLDS and RLDS does not apply.
- The SLDS and RLDS chronologically stores each change to any single database record. If a record changes 100 times since the last backup of the data set, the SLDS or RLDS includes all 100 changes. Yet, during recovery, you are only interested in the value the data had at the moment the data set was lost; the other 99 changes are irrelevant.

You can use the Database Change Accumulation utility to sort through your accumulated SLDSs and RLDSs in advance and condense and streamline them. This utility:

- Picks out only those log records relating to recovery of databases
- Sorts these records by data set within a database
- Finds the most recent change in each part of an individual record

As the utility creates the change accumulation data set, IMS compresses repeated single characters, such as blanks and zeros. IMS expands the data again during recovery.

Running the Database Change Accumulation utility is not required, but using it periodically speeds database recovery. Alternatively, you can run the Database Change Accumulation utility only when the need for recovery arises (just before running the Database Recovery utility). Running these two utilities instead of just the Database Recovery utility can reduce the total time needed for recovery, depending on how much unaccumulated log information exists.

**Related reference:**

 Database Change Accumulation utility (DFSUCUM0) (Database Utilities)

*Input to the Database Change Accumulation utility:*

In addition to using archived log data (SLDS and RLDS) as input to the Database Change Accumulation utility, you can also use a subset of the IMS log or a previous change accumulation data set. The utility writes the accumulated changes to a new change accumulation data set.

If the log data is on tape, you can specify all log volumes or a subset of log volumes as input to the Database Change Accumulation utility. When you specify a subset of log volumes, DBRC checks whether the subset is complete for each DBDS. A subset of log volumes is complete for a DBDS when all of the following conditions are true:

- The first volume in the subset is the volume with the first change to the DBDS since any of the following events occurred:
  - The last change accumulation.
  - The last image copy.
  - DBRC created the ALLOC record for this area (if the image copy was concurrent).
  - Checkpoint IDs that are found on logs prior to the run time of the image copy are included to ensure all updates are considered if the image copy was taken while updates were being made.
- The remaining volumes are in sequence.
- In a data-sharing environment, all logs containing changes for a DBDS are included.

Use the DBRC GENJCL.CA command to specify the subset of log volumes. You can request a specific number of log volumes either by volume (use the VOLNUM keyword) or by time stamp (use the CATIME keyword).

You can use a change accumulation data set as input to a later run of the Database Change Accumulation utility whether your subset of log volumes is complete or incomplete; however, you can use a change accumulation data set as input to the Database Recovery utility only if it represents a complete log subset.

**Related reference:**

 GENJCL.CA command (Commands)

*Purge time for database recovery and Database Change Accumulation utility:*

A purge time is the timestamp DBRC uses to start collecting changes for a database or area that might need to be recovered using either an image copy or a

HALDB Online Reorganization as a starting point for a recovery. Changes prior to this timestamp are purged from change accumulation data sets and are not considered for database recovery.

### Image copy purge times

If updates are not occurring (no active allocation records in RECON for the database), the run time of the image copy is used as the purge time. If updates are occurring when an image copy is run, DBRC checks checkpoint timestamps on logs previous to the run time of the image copy. For full function databases (including HALDBs), only one checkpoint on a previous log is considered. For Fast Path databases, two checkpoints on previous logs are considered. In order to keep purge times moving forward in time, ensure checkpoints exist on each log for each IMS updating the database.

The run time of the image copy is selected as the purge time for the following types of image copies for both full function and Fast Path databases:

- Database Image Copy (DFSUDMP0) utility was used to create the image copy while the database was unavailable for update (BATCH).
- Database Image Copy 2 utility invoked DFSMS Fast Replication to take an image copy while the database was unavailable for update processing (SMSOFFLC).
- Database Image Copy 2 was used to create the image copy while the database was unavailable for update processing (SMSNOCIC).
- Online Database Image Copy utility was used to create the image copy (ONLINE).

For other types of fuzzy image copies, DBRC uses an earlier checkpoint ID to ensure that all the changes are included for the Database Recovery or the Database Change Accumulation utility, as long as the database is allocated when the fuzzy image copy is run. If the database is not allocated by any IMS system when the fuzzy image copy is run, the run time of the image copy is selected as the purge time for the database.

Active allocations on all IMS systems when a fuzzy image copy is taken are considered and the earliest overall purge time is selected as follows when an active allocation is found for an IMS system:

- Database Image Copy utility was used to create the image copy while the database was available for update processing (CONCUR). A concurrent image copy is a “fuzzy” copy that occurs when updates are also being done. The data set uses logs in order to complete the image.
- Database Image Copy 2 utility invoked DFSMS Fast Replication to take an image copy concurrently with update processing (SMSONLC). The image is copy is a “fuzzy” copy so logs must be applied to recover the data set to a usable state.
- Database Image Copy 2 was used to create the image copy while the database was available for update processing (SMSCIC). The image copy is a “fuzzy” copy so logs must be applied to recover the data set to a usable state.

DBRC uses the checkpoint ID on logs to determine purge times with fuzzy image copies to ensure all updates are included as follows:

- **For full function databases (including HALDBs):** A checkpoint ID found on a log volume previous to the start time of the image copy is selected if the checkpoint ID is earlier than the allocation. Otherwise, the allocation time is selected as the purge time.



- **For Fast Path databases:** When at least two checkpoints are found on log volumes previous to the start time of the image copy, the log volume start time is used if the start time is earlier than the allocation. Otherwise, the allocation time is selected as the purge time.

If a HALDB Online Reorganization (OLR) is found that can be used as input to recovery for a HALDB that is later than the purge time of the image copy selected, the run time of the OLR is used as the purge time.

#### *Change accumulation groups:*

You can use DBRC to group DBDSs for which the Database Change Accumulation utility accumulates changes. These groups of DBDSs are called change accumulation groups.

You can define a change accumulation group by using the DBRC INIT.CAGRP command to write a CAGRP record in the RECON data set. Within this record, DBRC lists the DBDSs that make up the change accumulation group, identified by their database names and data set DD names. A change accumulation group can have a maximum of 32,767 members.

Before you can use the INIT.CAGRP command to define the change accumulation group, you must identify each member to DBRC using an INIT.DBDS command. A DBDS can belong to only one change accumulation group.

You can add or delete members of a change accumulation group using the CHANGE.CAGRP command.

**Restriction:** Do not issue a CHANGE.CAGRP command while the Database Change Accumulation utility is running because you could damage your database integrity.

#### *Defining change accumulation data sets for future use:*

You can define change accumulation data sets for future use for a given change accumulation group. Use the DBRC INIT.CA command to inform DBRC that these data sets exist, and specify the REUSE keyword on the INIT.CAGRP command when you define the group.

The GRPMAX keyword of the INIT.CAGRP command determines how many change accumulation data sets you can define.

#### *Reusing change accumulation data sets:*

To allow DBRC to reuse old change accumulation data sets, define a change accumulation group with the REUSE keyword and use the DBRC GENJCL.CA command to generate the JCL for the Database Change Accumulation utility job.

The Database Change Accumulation utility reuses the oldest change accumulation data set when all available change accumulation data sets for a particular change accumulation group have been used and the maximum number of change accumulation data sets has been reached. Reusing a change accumulation data set means that DBRC uses its data set name, volumes, physical space, and record in the RECON data set as if they were for an empty change accumulation data set.



If you define a group with the NOREUSE keyword, rather than reuse the data set, DBRC deletes the RECON record for the oldest change accumulation data set. In this case, DBRC does not scratch the data set. You must scratch the data set or keep track of it, because DBRC is no longer aware of it.

### **Specifying your choices for the log data sets:**

These topics describe the choices you must make when defining the various data sets involved in logging.

**Related reading:** For information on defining CQS data sets, see *IMS Version 11 System Definition*.

*Defining online log data sets:*

These topics describe tasks related to defining OLDSs.

*Choosing dual or single OLDS logging:*

IMS can log information to a single data set or to two identical data sets.

**Definitions:** *Single logging* uses a single data set, *dual logging* uses two data sets, where both data sets are identical. Whether you use single or dual logging, IMS writes information to sets of data sets, as described in “Online log data set” on page 94.

**Recommendation:** Use dual logging whenever possible because the OLDS is of primary importance to system integrity. Specify dual logging in IMS.PROCLIB using the OLDSDEF statement.

With dual logging, IMS has two options when one of the OLDSs (of a pair) gets an I/O error. The first option is to discard the pair and switch to a good pair. This behavior is just like single logging mode.

The second option is to not have IMS discard the pair of OLDSs unless both pairs of OLDSs fail. Then, if all the good pairs of OLDSs fail, IMS can degrade to single logging mode and use the good OLDS from each pair. From this point, IMS behaves just like if it had been in single logging mode from the beginning. This option also simplifies operational procedures because, for example, you only need to run the Log Recovery utility to clean up the OLDSs when you use (single or dual) logging and a write error occurs.

Essentially, there are three logging states for IMS to run in:

- Single logging, which has no redundancy and has the highest maintenance when an error occurs. When only two good OLDSs remain, IMS will terminate.
- Dual logging with DEGRADE=NO, which behaves just like single logging except that data is written to two LOGS instead of one. And, like in single logging mode, if only two good pairs of OLDSs remain, IMS will terminate.
- Dual logging with DEGRADE=YES. In this case, when each pair of OLDSs has at least one write error, IMS will switch to single logging mode and start logging to whichever OLDSs in the remaining pairs are good.

In all three cases, when IMS gets to the point where it only has two data sets left to write to, it terminates. If, for some reason, IMS has no good data sets to write

to, it will terminate with an ABENDU0616. For more information about log errors, see *IMS Version 11 Operations and Automation*.

#### *Defining the number of OLDSs:*

You must define at least three OLDSs (or OLDS pairs) to start IMS. However, you can define additional OLDSs (up to 100).

You must define the OLDSs to be used during initialization in the IMS.PROCLIB data set using the OLDSDEF statement. You can then dynamically allocate additional OLDSs. Specify the OLDSs you want to dynamically allocate using the DFSMDA macro; later you can use the /START OLDS command to add an OLDS.

**Related reading:** For more information on the DFSMDA macro, see *IMS Version 11 System Definition*.

When deciding how many OLDSs to define, consider the frequency of archiving and the amount of data you want to keep online. IMS reuses an OLDS only after IMS archives it. The number of OLDSs you define should be consistent with the frequency of archiving; if, for example, you archive frequently, you can probably plan on defining fewer OLDSs.

**Recommendation:** To avoid system failure, you should define more than the minimum number of OLDSs, even when using dual logging.

#### *Defining space for each OLDS:*

When defining the size of each OLDS, you also need to consider the size and location (DASD or tape) of the SLDS. If the SLDS is on tape, consider the size of an SLDS volume and how often you intend to archive the OLDS.

If the SLDS is on DASD, you should allocate enough space to contain all of the OLDSs to be archived. You might want to assign enough space to each OLDS so it fills an SLDS volume when it is archived. Or you might want to make each OLDS half the size of an SLDS volume, so you can initiate archive when two OLDSs are full.

Also, when defining the size of each OLDS, consider the amount of data you want to have online for doing such tasks as emergency restart or restarting a BMP. For example, if your system has a lot of activity and processes many transactions, consider defining larger OLDSs; then the records necessary for an emergency restart will more likely be online.

#### *Defining a block size for the OLDS:*

The main factor that determines OLDS block size is the track size of the OLDS devices. The OLDS block size cannot exceed the device track size. Define a block size that maximizes the amount of log data per track (for example, half-track for 3380 or 3390 DASD). Because IMS writes only full OLDS buffers to the OLDS, a large OLDS block size results in more efficient use of DASD space.

You must choose the block size for the OLDS carefully because changing the size of the OLDS after it has been established requires that you stop online work, archive all OLDSs, and scratch and reallocate them to make sure that their block sizes remain identical. When scratching and reallocating, you must delete OLDS entries from the DBRC RECON data set. If you are changing the OLDS block size

from being a multiple of 2048 (not running in z/Architecture® mode) to being a multiple of 4096 (running in z/Architecture mode) to ensure that the WADS gets reformatted according to the change in the OLDS block size, you must run the /NRE FORMAT WA or /NRE FORMAT ALL command. After you change the block size of an OLDS, you can restart online work only from an SLDS.

**Recommendation:** Take a checkpoint soon after the restart so later restarts can use the new OLDS, rather than the SLDS.

The block size of each OLDS must be the same. The OLDS block size must meet all of the following requirements:

- The block size must be a multiple of 2048 bytes (2 KB) if IMS is not running in z/Architecture mode. The block size must be a multiple of 4096 (4 KB), if IMS is running in z/Architecture mode.
- The block size must be at least 6144.

**Recommendation:** Use an 8 KB minimum or use the length of the largest message segment.

- The block size must not exceed a maximum of 30,720 bytes. This is the largest multiple of 2048 supported by BSAM.
- With the support of extended format OLDS, the calculation of how many blocks fit on a track requires that you add an extra 32 bytes of SAM overhead to the IMS block size before making the calculation.

As of IMS Version 10, your OLDS can be more than 65,535 tracks if you use the large sequential data set support. To take advantage of this support, hardware that has more than 65,535 tracks must be used.

During initialization, IMS ensures that the block size specified for the OLDS is large enough to handle the maximum length log record. If the block size specified is too small, IMS discards the OLDS data set looks at the next OLDS. If, at the end of initialization, there are not at least three pairs of usable OLDSS, IMS terminates with an 0073 abend.

#### *Defining devices for the OLDS:*

You configure the OLDS so that a system failure that renders an OLDS inaccessible does not stop the entire system. If you define dual OLDS logging, define each data set in an OLDS pair on different devices and, if possible, on different control units and channels.

The following figure shows an OLDS configuration with multiple devices. In this configuration, if one device fails, you can still access the OLDSS on the other two devices. For example, if device B fails while logging to primary OLDS 1 and secondary OLDS 1, you can still use primary OLDS 1 as input to the Log Archive utility, and IMS continues logging using the next available pair of OLDSS: Primary OLDS 3 and Secondary OLDS 3. If both Primary OLDS 3 and Secondary OLDS 3 have not been archived, IMS continues logging with the next available pair, Primary OLDS 6 and Secondary OLDS 6. The primary OLDS steps from one device to the next and the sequence wraps around.




Figure 24. Sample OLDS configuration

#### Changing OLDS characteristics:

Before you scratch and reallocate an OLDS that has been archived, you must delete the log control record in the DBRC RECON data set for this OLDS by using the DBRC DELETE.LOG OLDS(ddname) SSID (name) command.

You must delete the log control record for an OLDS that has been scratched and reallocated because the log control record only indicates that the OLDS has been archived. If you need this OLDS for IMS restart or batch backout, DBRC indicates to IMS to use this OLDS instead of the SLDS created by the archive job.

#### Related concepts:

Chapter 29, "Overview of DBRC," on page 479

#### Defining OLDS buffers:

You can define from 2 to 9999 OLDS buffers (the default is 5). The number of buffers can be changed during IMS restart.

If you change the number of buffers, each additional buffer might require additional tracks in the WADS. To calculate how many WADS tracks are required for each additional buffer, use the following formula:

$$(\text{OLDS block size}/2048) + 1$$

You might want to increase the number of OLDS buffers in the following circumstances:

- If IMS frequently waits for OLDS buffers
- If you have a high frequency of dynamic backout

Specify the number of OLDS buffers in IMS.PROCLIB using the OLDSDEF statement.

### *Defining system log data sets:*

To define system log data sets (SLDSs), consider choosing dual or single SLDS logging and defining a block size for the SLDS.

- Choosing dual or single SLDS logging

You can choose either single or dual logging for the SLDS, just as you can for the OLDS and WADS. When archiving to tape, you can force the primary and secondary volumes to contain the same data by specifying the number of log blocks per volume using the force-end-of-volume (FEOV) keyword of the SLDS control statement. When IMS writes the specified number of blocks, IMS forces end of volume on both the primary and secondary data sets.

To use dual logging, supply both a primary and secondary DD statement for each SLDS.

- Defining a block size for the SLDS

The block size of the SLDS can differ from the block size of the OLDSs being archived. However, the block size of primary and secondary SLDS must be the same when using the FEOV keyword.

### *Specifying the recovery log data set:*

To use dual logging for the RLDS, supply both a primary and secondary DD statement for each RLDS in the JCL for the Log Archive utility.

### *Specifying the restart data set:*

Allocate one cylinder of space for the restart data set (RDS). Although the checkpoint ID table uses one track, IMS also writes other recovery information to the RDS.

## **Contents of the log:**

IMS records activity on the OLDS. CQS records activity on a z/OS log stream. Each different activity is recorded as a separate log record.

Generally, you do not need to know the content of log records. IMS and CQS identify the correct records and use them automatically when they perform recovery.

The following items, however, should be kept in mind with respect to the logs:

- Log reduction

As part of log reduction, IMS compresses log data. IMS compresses repeated single characters, such as blanks and zeros, in the segment data portions of log records. This compression applies to updates resulting from DL/I insert (ISRT), delete (DLET), and replace (REPL) calls.

The counterpart to compression is expansion. During backout and database recovery, IMS expands the data in the database buffer. For change accumulation data sets, IMS expands the data in the change accumulation record.

- Logging and the Data Capture exit routine

IMS does not write log records for the Data Capture exit routine to show that it has been called, nor does IMS write the exit routine name in any log records. IMS does not differentiate between the application program and the exit routine. If you use the Data Capture exit routine extensively, your IMS system accounting and performance monitoring information is likely to show more system use for the application programs than they actually use.

**Related reading:** If you need to examine log records to solve a complex recovery problem, see *IMS Version 11 Diagnosis*.

## Backup

A *backup copy* of a data set serves the same purpose for a recoverable system as a checkpoint: it defines a place from which you can restart processing.

IMS provides utilities to allow you to make several types of backup copies.

### Database backup copies:

When IMS takes a regular system checkpoint, it records internal control information for DL/I and for Fast Path, but it does not record any of the contents of the database. If the database is lost, examining the last system checkpoint does not allow you to recover it. The system log can tell you what changes have occurred, but without the original database itself, recovery can be impossible.

**Recommendation:** Make a backup copy of all databases after you initially load them. You should also make new backup copies at regular intervals. More recent backup copies require fewer log change records to be processed during recovery, and thus the time needed for recovery is reduced.

IMS provides several utilities for making backup copies of a database:

- Image copy utilities

The IMS database image copy utilities allow you to take a “snapshot” of a database before and after changes have been to the database. These snapshots are called *image copies*.

**Definition:** An image copy is an as-is image of a database. The image copy utilities do not alter the physical format of the database as they copy it. Image copies are backup copies of your data that help speed up the processes of database recovery and backout.

The image copy utilities are Database Image Copy utility (DFSUDMP0), Database Image Copy 2 utility (DFSUDMT0), and Online Database Image Copy utility (DFSUICP0).

- The HISAM Reorganization Unload utility (DFSURUL0)

This utility allows you to process an entire HISAM database in one pass (the image copy utilities process each database data set individually) while it is unloaded and reorganized. This utility runs while the database is offline.

**Recommendation:** Make an image copy of the database before unloading it. You can use the image copy for recovery purposes if a failure occurs in the unloading process. You can also create an equivalent of an image copy data set without using the HISAM Reorganization Reload utility.

After you unload the database, you use the HISAM Reorganization Reload utility (DFSURRL0) before bringing the database back online. If you do not reload the database, an application program can use the database, but will use the old organization of the data set. Your backup copy will not match the log records produced for the database, and you will not be able to use the backup copy to recover the database without damaging your data integrity.

You can run these utilities with or without DBRC. You can also use various z/OS utilities to make your backup copies, but these utilities do not interact with DBRC, which could cause integrity problems depending on how your IMS system is defined.

**Related concepts:**

“Concurrent image copy” on page 519

**Related reference:**

 HISAM Reorganization Reload utility (DFSURRL0) (Database Utilities)

**Message queue backup copies:**

Messages (the transaction requests entered by end users, and the responses going back to them) are stored on queues before being processed. If you are not sharing the IMS message queues, messages are stored in the message queue data set, which resides partly on disk and partly in virtual storage.

In a shared-queues environment, IMS messages are kept on a coupling facility. When IMS takes a regular system checkpoint, it does not record the contents of the message queues, just as it does not record the contents of the databases.

When you shut down IMS normally (rather than abnormally), any changed messages in virtual storage (not on a coupling facility) are automatically written to the disk portion of the message queue data set. Therefore, if you periodically shut down IMS, there is little or no need to backup the message queues manually.

When you run IMS for extended periods without shutting it down, you might want to back up the message queues periodically. In a non-shared-queues environment, use the /CHECKPOINT SNAPQ command to back up the message queues; this command does not shut down IMS.

In a shared-queues environment, you must periodically back up the shared queues. Use the /CQCHKPT command to copy IMS messages to the Common Queue Server's structure recovery data set.

Backing up the message queues reduces the time required for recovery if problems arise with the message queue data sets.

**System data set backup copies:**

Database data sets and message queue data sets are not the only data sets for which backup copies should be made. You should also make backup copies of the IMS system data sets. These include the ACB library (IMS.ACBLIBx) and the MFS library (IMS.FORMATx).

IMS does not provide any special utilities or commands to make backup copies of system data sets. You can, however, make periodic backup copies using z/OS utilities, such as IEBCOPY. You can make these copies at the same time you make periodic backup copies of z/OS system libraries. You may also find it convenient to back up the system data sets and databases at the same time.

When you make online changes, you should make backup copies of the active data sets after you switch the inactive and active data sets.

You should also periodically back up RECON data sets using the BACKUP.RECON command.

**Recovery utilities and services**

IMS has a number of utilities and services to help you recover and maintain a recoverable system.



These utilities and services are:

- Database-related utilities: Use these utilities to make image copies of your databases, to accumulate and sort records of database changes, to speed up recovery, and to recover databases.
- System-related utilities: Use these utilities to help you manage the system by managing and recovering the log data sets or generating analysis reports.
- Transaction-Manager-related utilities: Use these utilities to help you manage and control your transactions.

IBM also offers a number of IMS productivity tools that can enhance your IMS environment. These tools can help you automate and speed up your IMS utility operations. They can also assist you in analyzing, managing, recovering, and repairing your IMS databases. You can learn more about these tools on the web at <http://www.ibm.com/software/data/db2imstools/products/ims-tools.html>.

## Database Recovery Control and recovery

The IMS Database Recovery Control (DBRC) facility makes it easier for you to recover IMS databases by extending the capabilities of IMS utilities for database recovery. DBRC can help recover both DL/I databases and Fast Path data entry databases (DEDBs).

DBRC offers two levels of control: log control and share control.

When you use log control only, DBRC controls the use and reuse of OLDSs for IMS.

When you use share control, DBRC:

- Controls the use and reuse of OLDSs for IMS
- Records recovery-related information in the Recovery Control (RECON) data set
- Assists in recovering databases
- Generates job control language (JCL) for recovery-related utilities
- Registers and controls the scheduling of databases

### Related concepts:

Chapter 29, “Overview of DBRC,” on page 479

## Automated operator interface

The IMS automated operator interface (AOI) allows application programs to issue IMS commands and can intercept IMS messages routed to the MTO.

You can use it to develop procedures tailored to your installation and to automate some parts of the recovery process. This facility is described in “Tools for Automated Operations” in *IMS Version 11 Operations and Automation*.

In an IMS complex with the Extended Recovery Facility (XRF), AOI runs on the active IMS subsystem, and after an XRF takeover, AOI runs on the new active subsystem.

## Recovery in an IMS DBCTL environment

The IMS DBCTL environment consists of the IMS Database Manager (not the Transaction Manager) and allows one or more transaction management subsystems, such as CICS, to access DL/I databases and DEDBs.

These transaction management subsystems are known as *coordinator controllers* (CCTLs). A CCTL communicates with the DBCTL subsystem using an interface



called the database resource adapter (DRA). Connections or paths between a DBCTL control region and a CCTL are created as part of CCTL initialization. These connections are called *threads*.

Because the CCTL is a separate subsystem, failures are isolated between the IMS DBCTL subsystem and the CCTL. That is, failure of the DBCTL subsystem does not generally cause the CCTL to terminate. Likewise, failure of a CCTL does not generally cause failure of the DBCTL subsystem. Generally, the only time a CCTL failure can affect the DBCTL subsystem is if the CCTL had one or more threads executing in DL/I at the time of its failure).

An IMS DBCTL subsystem cannot restart from log data not created by a DBCTL subsystem. Likewise, non-DBCTL IMS subsystems cannot restart using log data created by a DBCTL subsystem. In other words, restarts of an IMS DBCTL subsystem must use log data produced by that system. You can, however, perform database recovery using a combination of log data sets created by other subsystems: IMS DB/DC, batch, or DBCTL.

Because DBCTL requires the DBRC log control facility, you must use DBRC with DBCTL. Log records produced by an IMS DBCTL system are compatible with an IMS DB/DC system (except for subsystem restart).

### **Recovery in an XRF complex**

Use the Extended Recovery Facility (XRF) as an alternate IMS subsystem that monitors the log data of the active IMS subsystem and takes over the processing load of the active subsystem if it experiences a failure.

Everything about non-XRF recovery mechanisms is also true in an XRF complex, but with XRF, IMS can do more to reduce the time that data is unavailable to users after an abnormal event occurs.

You can specify conditions that trigger an automatic takeover, such as:

- A failure of the IMS control region
- A total failure of z/OS
- A single central processor complex (CPC) failure
- An internal resource lock manager (IRLM) failure that requires an IMS STATUS exit routine
- A Virtual Telecommunications Access Method (VTAM) failure that requires an IMS TPEND exit routine

You can also manually initiate an XRF takeover in a non-failure situation to introduce planned changes to your system with minimal disruption to users.

#### **Related concepts:**

Chapter 44, “Extended Recovery Facility Overview,” on page 585

### **Recovery with a Remote Site Recovery system**

Remote Site Recovery (RSR) is a separately-priced feature of IMS. It relies on having another IMS subsystem, situated on another z/OS system, that tracks the update activity on the primary IMS subsystem (or subsystems) to provide a backup.

RSR can track details of IMS full-function databases and Fast Path DEDBs at an alternate site. The remote site is connected to the site with the active systems by a

network connection using the VTAM APPC protocol. The VTAM connection is between separate IMS Transport Manager Subsystems (TMSs) on the active and tracking machines.

The IMS logger component on the active machines send the log data from all active IMS systems (DB/DC, DCCTL, DBCTL and batch) that are defined for RSR tracking to the tracking machine.

The TMS on the tracking machine receives this data and passes it to a single IMS region. This region processes the data and records it using normal IMS processing. Depending on what level of tracking has been requested, the IMS region might also apply the updates to the IMS databases.

If any interruptions to the network connection occur, RSR notes the gaps in the logging, obtains the missing log data, and performs catch-up processing when the link is reestablished.

The IMS system on the tracking machine normally can process only input from the TMS. The tracking machine only becomes a fully functioning system if it has to take over.

Not all databases are tracked. You define the databases that are tracked when you register them to DBRC.

**Related concepts:**

“RSR overview” on page 663

## Backward recovery

When IMS or an application program fails, you need to remove incorrect or unwanted changes from the database. Backward recovery or *backout* allows you to remove these incorrect updates.

The three types of backout are:

- Dynamic backout
- Backout during emergency restart
- Batch backout

IMS performs the first two types of backout automatically, and you initiate the last one manually.

IMS automatically (dynamically) backs out database changes in an online environment when any of the following occurs:

- An application program terminates abnormally
- An application program issues a rollback call
- An application program tries to access an unavailable database
- A deadlock occurs

In a batch environment, you can specify that IMS should dynamically back out database changes if the batch job abends, or issues a rollback call.

During restart processing after a system failure, IMS determines if any application programs were executing at the time of failure and if they made changes that need to be backed out. Before IMS restarts, it scans log records for these changes, and

then backs out the changes from the affected databases. If IMS runs out of memory while scanning the log, you could get a message telling you to back out the changes manually.

You can use the IMS Batch Backout utility (DFSBB00) to remove database changes made during execution of a DL/I or DBB region or an online program. You can use the utility to back out changes since the last checkpoint. You can select a specific checkpoint if the batch region does not use data sharing, and if it is not a BMP. For BMPs, do not specify a checkpoint because the utility always backs out BMPs to the last checkpoint on the log.

If dynamic backout or backout during emergency restart fails, IMS stops the databases for which backout did not complete, and retries the backouts when you restart the databases.

## Forward recovery

Forward recovery involves reconstructing information from the IMS logs and re-applying it to a database.

With forward recovery, you reconstruct information or work that has been lost, and add it to the database. With backout, you remove incorrect or unwanted changes from information. For an IMS subsystem, you perform the task of forward recovery (IMS does supply utilities to help). Backout of transactions (and their associated database updates) is handled automatically by IMS. Backout usually does not require your involvement, but an application program can control backout processing for its own transactions. Backout of batch programs can be done either automatically or using an IMS utility.

## Shutdown of the IMS system

Certain problems are severe enough that they cause the IMS subsystem to fail. In these cases, IMS shuts itself down. In other cases, however, IMS keeps running, so you must shut it down manually before you can perform recovery. You can shutdown either part of the IMS subsystem (partial shutdown), or you can shut down all of it (full shutdown).

**Recommendation:** Because a primary goal is to keep IMS function available to users, you should only use a full shutdown when a partial shutdown does not recover the system.

When only a part of IMS is malfunctioning, you might be able to shut down only that part and leave the rest of IMS functioning productively. For example, a faulty database can be taken offline (made unavailable to application programs), while IMS and the other databases continue processing. Or a terminal that is malfunctioning can be detached, while all other functions continue unaffected.

To shut down only part of IMS, use a command that stops the component that is malfunctioning.

You can also shut down IMS in a controlled manner. A controlled shutdown is desirable because it saves current information in the system, and allows an easy and accurate restart of the system at a later time.

To shut down IMS, use the /CHECKPOINT command with one of the following keywords: FREEZE, DUMPQ, or PURGE.

**Related reference:**

## Recovery during restart

IMS can recover from many types of failures during restart. In some cases, the recovery is automatic, for example, backout of transactions and the associated database updates. In other cases, you request a type of recovery on the restart command, for example, you can tell IMS to rebuild the message queues from a previous backup copy.

Generally, the way you restart IMS corresponds to the way you shut it down, or the way it fails.

- If you only shut down part of IMS (for example, take a database offline or detach a line), you need to restart only that part.
- If you shut down all of IMS or if IMS fails, you need to restart the whole IMS subsystem.

After recovering a failed component, you can restart it. For example, if you take a database offline because of an I/O error and then recover it, you can again make it available to applications.

To restart only part of IMS, use a command that starts the component that is stopped.

Before you can restart the entire IMS subsystem, you must first perform all necessary recovery. You can then restart IMS in one of three ways:

- *Normal restart* initializes an IMS subsystem that has not failed. A normal restart can be either a *cold start* which restarts IMS without reference to any previous execution of IMS, or a *warm start* which restarts an IMS subsystem that was terminated using a /CHECKPOINT command. The normal restart command is /NRESTART.
- *Emergency restart* initializes an IMS subsystem that has failed. During an emergency restart, IMS resets transactions and active regions, and restores databases and message queues to the most recent sync point. You must manually restart batch and batch message processing (BMP) regions. The emergency restart command is /ERESTART.
- *Automatic restart* reduces MTO intervention and can make restart faster because IMS automatically chooses the appropriate restart command. The operator does not enter a restart command; instead you specify automatic restart by coding AUTO=Y in the JCL for the IMS control region.

## Complications in recovery

If you could always recover a system simply by knowing where you started and keeping track of what you have done since then, it would be easy. Unfortunately, complications often arise, not only in the data and applications, but in the mechanisms of recovery themselves.

The IMS logs (on DASD or tape) are just as subject to I/O problems as any other data set. You can run out of space, or develop physical problems with the drive. The internal IMS modules involved in recovery are totally dependent on the operating system and hardware; if the system crashes for any reason (for example, power failure, abend, hardware malfunction), the recovery processing performed by these modules might be incomplete.

So, recovering from a problem can be a two-part process: first you must recover the IMS recovery tools themselves, then you can recover the data and applications.

## What IMS cannot recover

IMS can automatically recover from some errors, and IMS provides mechanisms to allow you to recover from many more errors, but there are some errors IMS cannot recover.

Generally, these unrecoverable errors are of two types:

- Application logic or input errors

When you detect an error caused by an application program that is running, backing out the faulty data is fairly simple. But if that program is given the wrong input or is constructed with logic errors, you might not be able to detect such problems because the program will run to completion. After the application program has finished running (or has committed its data), the faulty data can be used by other programs, and the problem can spread throughout your system.

IMS has no automatic or guaranteed way to back out committed data. You might be able to back out the data manually (for example, by editing individual records in the affected databases), but such a process is likely to be difficult. And the more time that passes after the original errors are introduced, the more programs are likely to have used the bad data, and the less likely your chances of success. Moreover, output that has been produced or actions taken as a result of the faulty data will already be outside the boundaries of the databases.

**Recommendation:** To minimize application logic or input errors, you should extensively test new applications (both individually and in an integrated system test) before bringing them online. Extensive validation of user input by application programs also helps minimize input errors.

- Operational errors, including misuse (accidental or malicious) of the IMS recovery facilities

Successful recovery requires both proper operation of IMS on a daily basis and proper use of the recovery facilities. This proper use includes such things as:

- Using the correct data sets
- Avoiding improper job cancellations during regular operations
- Regularly maintaining the logs
- Using the correct log volumes in the correct sequence
- Using the appropriate utilities at the appropriate time

You should have no difficulties in operating IMS properly, and other topics in this information explain how to set up the necessary procedures and guidelines. However, if you operate IMS, its recovery facilities, or both, improperly, IMS might not be able to correct problems introduced by such improper use.

## Overhead of recovery

Although the value of recovery is obvious, it is not without cost. Logging, for example, takes time and space: some amount of processing time is required when a log record is created, and each log record uses space on a data set. Each checkpoint written also takes time and space, as does each backup copy made.

Having a basic recovery scheme is a practical necessity, not really a matter of choice. However, you do have a choice about how elaborate to make your recovery scheme.

Generally, the more extensive recovery preparations you build into your daily operations, the faster you can bring your system back into production after an error occurs. The trade-off is between speed of recovery and the daily overhead of recovery maintenance. You also need to consider the cost of having your online system down.

In judging this trade-off, consider how frequently errors occur that require recovery. Errors will certainly occur, but rare errors merit less recovery preparation than those that occur more frequently. Your decision about the extent to which you prepare for recovery mechanisms should be based on knowing your installation's needs and priorities, and weighing them against each other.

---

## IMS Application Menu

The IMS Application Menu provides a common interface to IBM-supplied IMS applications that run on TSO using ISPF.

The following applications run on TSO using ISPF:


- Single Point of Control (SPOC)
- Manage Resources
- Knowledge-Based Log Analysis (KBLA)
- HALDB Partition Definition utility
- IMS Syntax Checker
- Installation Verification Program (IVP)
- IVP Export utility
- IPCS with IMS Dump Formatter (IPCS)
- Abend Search and Notification (ASN)

**Tip:** To provide access to the IMS Application Menu, include the IMS.SDFSEEXEC data set in the SYSPROC DD concatenation.

**Attention:** Ensure that IPCS is started before the IMS Application Menu is started. Otherwise, message DFSIX103 is displayed.

Use the DFSAPPL command to start the IMS Application Menu. You can either use a TSO command or an EXEC command:

- TSO %DFSAPPL HLQ(*myhlq*)
- EXEC 'IMS.SDFSEEXEC(DFSAPPL)' 'HLQ(*myhlq*)'



### Notes:

- 1 The **HLQ** parameter is required the first time you use the command. Thereafter, **HLQ** is an optional parameter.

Where:

## DFSAPPL

Command to start the IMS Application menu

**HLQ** Keyword that enables you to specify the high-level qualifier of the IMS distribution data sets

The HLQ parameter is required the first time you use the command. If you do not specify it, the command uses the most recently specified high-level qualifier. This parameter is optional.

*myhlq* High-level qualifier of the IMS distribution data sets

## ALTRESL

Keyword that enables you to specify a list of data set names that contain load modules

If you specify the ALTRESL parameter, you should include SDFSRESL in the list of data set names. If you do not specify the ALTRESL parameter, *myhlq.SDFSRESL* is used as the ISPLLIB data set.

*myhlq.data\_set\_name1*

Fully-qualified name of a data set that contains load modules

**Note:** Some applications require an ISPTABL data set. If the ISPTABL data set is allocated, it will continue to be used. If the ISPTABL data set is not in use, a new one is allocated using your TSO prefix or userid as the high-level qualifier.

The IMS Application menu is shown in the following figure.




Figure 25. IMS Application Menu

Using the IMS Application menu, you can start any of the TSO or ISPF applications by selecting the application and pressing the Enter key.

You can also link to the IMS Application menu from your local ISPF option menu. The following panel is an example:

```
)BODY
      Local Option Menu
Option ==>_ZCMD

      I    IMS Application Menu
      .
      .
)PROC
&ZSEL = TRANS(TRUNC(&ZCMD, '.'))
```

```
I,'CMD(%DFSAPPL HLQ(myhlq)) NEWAPPL(DFS) NOCHECK'  
.  
.  
)END
```

**Related reading:** For more information about any of the ISPF panels within the IMS Application Menu, see the ISPF online help that comes with the applications.

**Related concepts:**

“Analyzing log records using KBLA utilities” on page 833

 Controlling IMS with the TSO SPOC application (Operations and Automation)

 IMS Syntax Checker (System Definition)

 IMS installation verification program (IVP) overview (Installation)

 Using IPCS and the IMS offline dump formatter (Diagnosis)

**Related reference:**

 HALDB Partition Definition utility (%DFSHALDB) (Database Utilities)

**Related information:**

 IMS abend search and notification (Messages and Codes)





---

## **Part 2. IMS system administration considerations and tasks**

These topics describe information and required steps for day-to-day operation of the IMS system.



---

## Chapter 5. z/OS interface considerations

These topics describe information and required steps that you must consider while installing IMS and IRLM on z/OS.

**Important:** After the z/OS and VTAM interface steps are completed, you must IPL z/OS and specify either CLPA or MLPA=xx, or both. Also, note that IMS can run in 31-bit or 64-bit processing modes.

---

### z/OS interface considerations for IMS

There are many requirements that you must consider and required steps to ensure a complete and correct installation of IMS on z/OS. These topics describe these requirements and required actions.

#### Preventing installation problems

Be sure to take the following actions to prevent problems during the installation of IMS on z/OS.

- Use z/OS macro libraries for your IMS stage 2 definition. IMS runs only under z/OS .
- Include the libraries from which IMS is loaded and executed in the appropriate authorization table, so that the control region executes as an APF-authorized program. In z/OS, IMS runs as an authorized program.

**Related reading:** For information about APF authorization, see IEAAPFxx (authorized program facility list) in *z/OS MVS Initialization and Tuning Reference*.

- Use JOBLIB or STEPLIB DD statements instead of having the IMS.SDFSRESL in LNKLSLxx (those data sets concatenated to SYS1.LINKLIB). If IMS.SDFSRESL is in LNKLSLxx, it is possible for a different IMS release level (whose own IMS.SDFSRESL is not properly APF authorized) to load the modules from LNKLSLxx. The incompatible module release level can cause unpredictable results.
- Update the program properties table. The IMS control region operates as a job step task or as a system task. All control region execution is in supervisor state. See “Updating the IBM-supplied z/OS Program Properties Table” on page 126 for more information.

**Related reading:** For additional information on maintaining system integrity when running under z/OS, refer to *z/OS MVS Programming: Authorized Assembler Services Reference*.

#### Setting up JCL

The following list contains the requirements for setting up your z/OS JCL.

- The JOB or STEP libraries must be APF authorized for the control region. For the dependent region, PGMLIB does not need to be authorized and can be concatenated with SDFSRESL as STEPLIB.
- The EXEC statement must specify PGM=DFSMVRC0 for the control region.
- The following libraries must be APF authorized:
  - IMS.SDFSRESL
  - IMS.MODBLKSA and IMS.MODBLKSB
  - IMS.SDXRRESL

- IMS.SDFSJLIB
- The library into which your DB2 modules are loaded (DFSESL DD or a JOBLIB or STEPLIB)

**Related reading:** For more information on z/OS JCL, refer to the information on the system definition process in *IMS Version 11 System Definition*.

## Configuring JES for dependent regions

The EOM broadcast is always made when an IMS dependent region running as a started task ends, but the EOM broadcast may not be made, and therefore not processed, when a batch job ends.

For IMS dependent regions run as batch jobs, an EOM broadcast is not made until both the job and the initiator used by that job ends. Therefore, if a dependent region is run as a batch job and it encounters a storage problem as described, the EOM clean up service will not be invoked until the initiator ends.

In a JES2 environment, initiators that are used for dependent regions can be set to stop automatically when the job ends by implementing an IBM-defined exit.

Use JES2 exit 32 for Subsystem Interface (SSI) Job Selection to specify that the initiator is to end (drain) or to end and restart when the job ends.

1. Set bit 4 in the response byte to end and then restart the initiator when the job ends.
2. Set bit 5 to end the initiator when the job ends (for more details see z/OS V1R12.0 JES2 Installation Exits SA22-7534-12 ).

In a JES3 environment, initiators that are used for dependent regions can be set to stop automatically when the job ends by specifying either DYNAMIC or DEMAND for the UNALOPT option of the EXRESC parameter in the initiator's GROUP definition (see z/OS JES3 Initialization and Tuning Guide SA22-7549-11 for more details).

## Keeping some required nonstandard z/OS macros in their original libraries

The assembly of certain IMS modules requires z/OS macros not contained on the standard z/OS System Macro libraries. Because these requirements are subject to change due to IMS and z/OS maintenance, keep these macros in their original libraries, and use the JCL generated by IMS for SYS1.MODGEN (or SYS1.AMODGEN).

## Updating the IBM-supplied z/OS Program Properties Table

The following modules are predefined in the default Program Properties Table (PPT) that is shipped with z/OS: BPEINI00, CQSINIT0, DFSMVRC0, and DXRRML00.

If you do not modify the default z/OS PPT, these IMS modules are already included in the PPT, and you do not need to take any action. If you have removed the default entries for these modules, you must reinstate the entries, which will be described in the succeeding sections.

**Related reading:** For information on updating the PPT, see *z/OS MVS Initialization and Tuning Reference*.

## Related tasks:

 Adding entries to the z/OS Program Properties Table (System Definition)

### Adding the IMS entry to the z/OS Program Properties Table

An IMS online environment (DB/DC, DBCTL, DCCTL) requires this PPT entry. If you are only using IMS BATCH, this entry is not needed.

A sample of the required entry is shown below and may be found in the IMS.INSTALIB data set. Refer to “IVP jobs and tasks” in *IMS Version 11 Installation* for the correct entry titled “Update SCHEDxx -- PPT Entries”.

To make this entry, edit the SCHEDxx member of the SYS1.PARMLIB data set. Add the following entry to the SCHEDxx member:

|     |                   |                                    |    |
|-----|-------------------|------------------------------------|----|
| PPT | PGMNAME(DFSVMRC0) | /* IMS ONLINE CONTROL REGION       | */ |
|     | CANCEL            | /* PROGRAM NAME = DFSVMRC0         | */ |
|     | KEY(7)            | /* PROGRAM CAN BE CANCELLED        | */ |
|     | NOSWAP            | /* PROTECT KEY ASSIGNED IS 7       | */ |
|     | NOPRIV            | /* PROGRAM IS NOT-SWAPPABLE        | */ |
|     | SYST              | /* PROGRAM IS NOT PRIVILEGED       | */ |
|     | DSI               | /* PROGRAM IS A SYSTEM TASK        | */ |
|     | PASS              | /* DOES REQUIRE DATA SET INTEGRITY | */ |
|     | AFF(NONE)         | /* PASSWORD PROTECTION ACTIVE      | */ |
|     | NOPREF            | /* NO CPU AFFINITY                 | */ |
|     |                   | /* NO PREFERRED STORAGE FRAMES     | */ |

The PPT Entry for program DFSVMRC0 must specify NOSWAP as shown.

### Adding the CQS entry to the z/OS Program Properties Table

If you are using CQS, either the CQSINIT0 or the BPEINI00 z/OS PPT entry is required.

A sample of the CQSINIT0 entry is shown below and can be found in the IMS.INSTALIB data set. A sample of the BPEINI00 entry is shown under the heading “Adding the BPE entry to the z/OS Program Properties Table.” Refer to “IVP jobs and tasks” in *IMS Version 11 Installation* for the correct entry titled “Update SCHEDxx -- PPT Entries”.

To make this entry, edit the SCHEDxx member of the SYS1.PARMLIB data set. Add the following entry to the SCHEDxx member:

|     |                   |                                    |    |
|-----|-------------------|------------------------------------|----|
| PPT | PGMNAME(CQSINIT0) | /* CQS - COMMON QUEUE SERVER       | */ |
|     | CANCEL            | /* PROGRAM NAME = CQSINIT0         | */ |
|     | KEY(7)            | /* PROGRAM CAN BE CANCELLED        | */ |
|     | NOSWAP            | /* PROTECT KEY ASSIGNED IS 7       | */ |
|     | NOPRIV            | /* PROGRAM IS NOT-SWAPPABLE        | */ |
|     | SYST              | /* PROGRAM IS NOT PRIVILEGED       | */ |
|     | DSI               | /* PROGRAM IS A SYSTEM TASK        | */ |
|     | PASS              | /* DOES REQUIRE DATA SET INTEGRITY | */ |
|     | AFF(NONE)         | /* PASSWORD PROTECTION ACTIVE      | */ |
|     | NOPREF            | /* NO CPU AFFINITY                 | */ |
|     |                   | /* NO PREFERRED STORAGE FRAMES     | */ |

The PPT Entry for program CQSINIT0 must specify NOSWAP as shown.

### Adding the BPE entry to the z/OS Program Properties Table

The BPE program properties table entry, BPEINI00, is used by various IMS address spaces.

These address spaces include:

- Common Queue Server (CQS)
- Database Recovery Control (DBRC) - optional
- IMS Connect
- Open Database Manager (ODBM)
- Operations Manager (OM)
- Resource Manager (RM)
- Structured Call Interface (SCI)

If you are using any of these address spaces, you must have the BPEINI00 entry in your PPT.

To make this entry, edit the SCHEDxx member of the SYS1.PARMLIB data set. Add the following entry to the SCHEDxx member:

```

PPT PGMNAME(BPEINI00)      /* BPE - BASE PRIMITIVE ENVIRONMENT      */
                                /* PROGRAM NAME = BPEINI00                  */
                                /* PROGRAM CAN BE CANCELLED                  */
                                /* PROTECT KEY ASSIGNED IS 7                  */
                                /* PROGRAM IS NOT-SWAPPABLE                  */
                                /* PROGRAM IS NOT PRIVILEGED                  */
                                /* REQUIRES DATA SET INTEGRITY                  */
                                /* CANNOT BYPASS PASSWORD PROTECTION          */
                                /* PROGRAM IS A SYSTEM TASK                  */
                                /* NO CPU AFFINITY                      */
                                /* NO PREFERRED STORAGE FRAMES              */
                                /*

```

To make the SCHEDxx changes effective, take one of the following actions:

- Restart the z/OS system.
- Issue the z/OS SET SCH= command.

## Adding the IMS Connect entry to the z/OS Program Properties Table

IMS Connect uses the predefined z/OS Program Properties Table (PPT) entry BPEINI00, but HWSHWS00 is also supported.

The examples shown below use HWSHWS00 instead of BPEINI00. You can include both entries in the PPT, for example, a mixed-version IMS environment that contains IMS Connect. HWSHWS00 is not included in the default PPT, so you must add HWSHWS00 to the PPT if you want to use it.

**Recommendation:** Use BPEINI00 on your IMS Connect startup JCL and then you do not have to put HWSHWS00 in the PPT.

## Examples

For TCP/IP communications only, add the following entry in the z/OS PPT:

```

PPT PGMNAME(HWSHWS00)      /* PROGRAM NAME = HWSHWS00      */
                                /* PROGRAM CAN BE CANCELED      */
                                /* PROTECT KEY ASSIGNED IS 7      */
                                /* PROGRAM IS SWAPPABLE          */
                                /* PROGRAM IS NOT PRIVILEGED      */
                                /* REQUIRES DATA SET INTEGRITY      */
                                /* CANNOT BYPASS PASSWORD PROTECTION */
                                /* PROGRAM IS A SYSTEM TASK          */
                                /* NO CPU AFFINITY                      */
                                /* NO PREFERRED STORAGE FRAMES              */
                                /*

```

If you are using local option for client communications, either by itself or with TCP/IP communications, add the following entry in the z/OS PPT:

```
PPT PGMNAME(HWSHWS00) /* PROGRAM NAME = HWSHWS00 */
      CANCEL           /* PROGRAM CAN BE CANCELED */
      KEY(7)           /* PROTECT KEY ASSIGNED IS 7 */
      NOSWAP           /* PROGRAM IS NOT SWAPPABLE */
      NOPRIV           /* PROGRAM IS NOT PRIVILEGED */
      DSI              /* REQUIRES DATA SET INTEGRITY */
      PASS             /* CANNOT BYPASS PASSWORD PROTECTION */
      SYST             /* PROGRAM IS A SYSTEM TASK */
      AFF(NONE)        /* NO CPU AFFINITY */
      NOPREF           /* NO PREFERRED STORAGE FRAMES */
```

## Adding the IRLM entry to the z/OS Program Properties Table

If you are using IRLM, the z/OS PPT entry is required.

A sample of the required entry is shown below and can be found in the IMS.INSTALIB data set. Refer to “IVP jobs and tasks” in *IMS Version 11 Installation* for the correct entry titled “Update SCHEDxx -- PPT Entries”.

To make this entry, edit the SCHEDxx member of the SYS1.PARMLIB data set. Add the following entry to the SCHEDxx member:

```
PPT PGMNAME(DXRRLM00) /* IRLM - RESOURCE LOCK MANAGER */
      CANCEL           /* PROGRAM NAME = DXRRLM00 */
      KEY(7)           /* PROGRAM CAN BE CANCELED */
      NOSWAP           /* PROTECT KEY ASSIGNED IS 7 */
      NOPRIV           /* PROGRAM IS NOT-SWAPPABLE */
      SYST             /* PROGRAM IS NOT PRIVILEGED */
      DSI              /* PROGRAM IS A SYSTEM TASK */
      PASS             /* DOES REQUIRE DATA SET INTEGRITY */
      AFF(NONE)        /* PASSWORD PROTECTION ACTIVE */
      NOPREF           /* NO CPU AFFINITY */
      NOPREF           /* NO PREFERRED STORAGE FRAMES */
```

The PPT Entry for program DXRRLM00 must specify NOSWAP as shown.

## Installing required IMS links to z/OS

You must install these modules and procedures on your z/OS system.

Stage 2 of IMS system definition might make the following modifications:

- Binds the following modules into IMS.SDFSRESL:
  - Type 2 SVC routine
  - DBRC Type 4 SVC routine
  - CTC Channel-end Appendage routine (if the MSC with the CTC option is defined)
- Copies cataloged procedures into IMS.PROCLIB

The following table provides is an overview of the actions needed in order for your IMS system to run under z/OS.

*Table 5. Steps required to run under z/OS depending on the IMS environment*

| Action                                    | DB<br>batch<br>system | DBCTL<br>system | DB/DC<br>system | DCCTL<br>system |
|-------------------------------------------|-----------------------|-----------------|-----------------|-----------------|
| Bind the Type 2 SVC with the z/OS nucleus | Yes                   | Yes             | Yes             | Yes             |



Table 5. Steps required to run under z/OS depending on the IMS environment (continued)

| Action                                                                                                                                                                                                                                                   | DB batch system | DBCTL system | DB/DC system | DCCTL system |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|--------------|--------------|--------------|
| Load the Type 2 SVC from SYS1.NUCLEUS using one of the following methods: <ul style="list-style-type: none"> <li>• The Nucleus Module Loader facilities</li> <li>• A SYS1.IPLPARM member, NUCLSTxx</li> <li>• A SYS1.PARMLIB member, NUCLSTxx</li> </ul> | Yes             | Yes          | Yes          | Yes          |
| Bind the DBRC Type 4 SVC module into LPALIB (or, optionally, into an MLPA library)                                                                                                                                                                       | Yes             | Yes          | Yes          | Yes          |
| Bind the CTC channel-end appendage (if your system has MSC with the CTC option) into LPALIB (or, optionally, into an MLPA library)                                                                                                                       | No              | No           | Yes          | Yes          |

The following table shows the modules that are required by the z/OS interface. The table shows the module name in its distribution library (IMS.ADFSLOAD) and the load module name in its target library (IMS.SDFSRESL) after the module is bound.

Table 6. z/OS interface modules.

| IMS.ADFSLOAD          | IMS.SDFSRESL          | Description                            |
|-----------------------|-----------------------|----------------------------------------|
| DFSVC200 <sup>2</sup> | IGCiii <sup>2</sup>   | Type 2 SVC Vector routine <sup>3</sup> |
| DSP00MVS <sup>1</sup> | IGC00nnn <sup>1</sup> | DBRC Type 4 SVC routine <sup>3</sup>   |
| DFSCMC10              | IGG019zz <sup>1</sup> | CTC channel-end appendage <sup>3</sup> |

**Notes:**

**iii** Specifies the Type 2 SVC number

**nnn** Indicates the signed decimal Type 4 SVC number, for example, SVC 255 is 25E

**zz** Indicates the channel-end appendage number specified on the IMSCTF macro

<sup>1</sup> These modules must be bound with the RENT and REFR attributes.

<sup>2</sup> This module must be bound with the RENT, REFR, and SCTR Binder options. The modules are placed in SYS1.NUCLEUS.

<sup>3</sup> These modules are bound as part of the system definition process.

## IMS SVC modules

IMS uses a Type 2 supervisor call (SVC), in the range of 200-255, for batch, utility, DBCTL, DCCTL, and DB/DC IMS control program functions. IMS uses a Type 4 supervisor call (SVC), in the range of 200-255, for DBRC functions. Specify these routines in IMS system definition.

If you are installing different release levels of IMS in the same z/OS system, note that the Type 2 SVC and Type 4 SVC are downward compatible. The IMS Version 11 level of the SVC can be used by IMS Version 9 and IMS Version 10. However, the IMS Version 9 level cannot be used by IMS Version 10 or IMS Version 11, and the IMS Version 10 level cannot be used by IMS Version 11.

IMS system definition creates the SVC routines using the IMSCTF macro-defined user-specified numbers, or the IMS-provided default numbers. IMS system definition copies the load modules representing the SVC routine into IMS.SDFSRESL.

When you define the IMS and DBRC SVCs to z/OS, follow this format:

**Example:**

```
SVCPARM 254,REPLACE,TYPE(2)
SVCPARM 255,REPLACE,TYPE(4)
```

**Related reading:** Refer to *z/OS MVS Initialization and Tuning Reference* for information on defining SVCs to z/OS.

## Installing the type 2 SVC module

The IMS Type 2 SVC must be incorporated into the z/OS nucleus.

**Attention:** The SYS1.NUCLEUS must not have secondary extents. z/OS cannot recognize secondary extents.

To incorporate the IMS Type 2 SVC into the z/OS nucleus, perform one of the following tasks:

- Bind the Type 2 SVC with the z/OS nucleus.  
You can bind the Type 2 SVC with the z/OS nucleus using one of the two following steps.
  1. Invoking the Binder utility through a batch job
  2. Creating and then performing a RECEIVE and APPLY for an SMP/E USERMOD
- Load the Type 2 SVC from SYS1.NUCLEUS using the Nucleus Module Loader facilities.
  1. Create a Nucleus Module List (NML) that contains the list of IMS SVCs that you want loaded into the z/OS nucleus. IMS uses the IEANS001 NML.
  2. Assemble and bind the Type 2 SVC into SYS1.NUCLEUS.

This method is included as an example in the IVP materials.

- Load the Type 2 SVC from SYS1.NUCLEUS using a SYS1.IPLPARM member, NUCLSTxx.
  1. Bind the IMS SVCs from IMS.SDFSRESL into SYS1.NUCLEUS.

**Attention:** Determine, from the z/OS systems programmer, the appropriate NUCLSTxx member to use. Note that the LOADxx member and its associated NUCLSTxx member must both reside in SYS1.IPLPARM. If the 2 members are not in this library, IMS will enter a **Disabled Wait** state and the IPL process stops.
  2. Define an INCLUDE statement for the IMS SVC in the NUCLSTxx member of SYS1.IPLPARM.
- Load the Type 2 SVC from SYS1.NUCLEUS using a SYS1.PARMLIB member, NUCLSTxx.
  1. Bind the IMS SVCs from IMS.SDFSRESL into SYS1.NUCLEUS.

**Attention:** Determine, from the z/OS systems programmer, the appropriate NUCLSTxx member to use. Note that the LOADxx member and its associated NUCLSTxx member must both reside in SYS1.PARMLIB. If the 2 members are not in this library, IMS will enter a **Disabled Wait** state and the IPL process stops.

2. Define an INCLUDE statement for the IMS SVC in the NUCLSTxx member of SYS1.PARMLIB.

## Binding the channel-to-channel (CTC) channel-end appendage

If you define multiple systems coupling (MSC) with the CTC option, bind the CTC channel-end appendage named IGG019zz, where zz is the CTC appendage number.

## Installing the resource cleanup module

Although DFSMRCL0 is not required for IMS Version 9 or later, this module is still shipped to support users who point to DFSMRCL0 directly in the IMS library. DFSMRCL0 from IMS Version 9 or later can also be used to provide resource cleanup for IMS Version 8 and earlier. You must have DFSMRCL0 installed on any z/OS LPAR where you will run a pre-Version 9 IMS (either online or batch). The version of the DFSMRCL0 module you install must be at least the same as the highest pre-Version 9 IMS that you will be running on the z/OS LPAR.

In IMS Version 8 and earlier, you must install the IMS module DFSMRCL0 in the host z/OS system as a static resource cleanup module. You are required to bind DFSMRCL0 into SYS1.LPALIB or an MLPA library. The module name DFSMRCL0 must be added to the IEAVTRML CSECT of z/OS module IGC0001C in SYS1.LPALIB.

In IMS Version 9 and later, IMS uses a dynamic resource cleanup module (DFSMRC20). No user setup is required; you do not need to install the static resource cleanup module (DFSMRCL0) on the host z/OS system.

**Recommendation:** Do not uninstall DFSMRCL0 from releases of IMS earlier than IMS Version 9 until your migration to IMS Version 9 or later is complete, and there is no possibility that you will run an earlier release of IMS. DFSMRCL0 and the dynamic resource cleanup module (DFSMRC20) can coexist on the same system.

## Examples

- You plan to run IMS code from all of IMS Version 7, IMS Version 8, and IMS Version 9 on the same LPAR. You must install DFSMRCL0 as a static resource cleanup module to provide cleanup functions for the IMS Version 7 and 8 systems. You can use the DFSMRCL0 shipped with either IMS Version 8 or IMS Version 9.
- You plan to run IMS code from IMS Version 8, IMS Version 9, IMS Version 10, and IMS Version 11 on the same LPAR. You must install DFSMRCL0 as a static resource cleanup module to provide cleanup functions for the IMS Version 8 system. You can use the DFSMRCL0 shipped with any of IMS Version 8, IMS Version 9, IMS Version 10, or IMS Version 11.
- You plan to run IMS code from IMS Version 9 and IMS Version 11 on the same LPAR. You do not need to install DFSMRCL0.

## Uninstalling DFSMRCL0

After you have completely migrated to IMS Version 9 or later and there is no possibility of running an earlier release of IMS (both IMS control and IMS batch jobs), you can remove DFSMRCL0 from the host z/OS system.

To remove DFSMRCL0 from the host z/OS system, perform the following steps:

1. Remove the name DFSMRCL0 from the IEAVTRML CSECT of module IGC0001C in SYS1.LPALIB. Removing this name prevents the operating system from installing DFSMRCL0 as a Static Resource Cleanup routine at the next IPL.
2. Remove module DFSMRCL0 from SYS1.LPALIB or the MLPA library where DFSMRCL0 was bound.
3. Restart with CLPA to enable these changes.

**Important:** You must perform these tasks in the order specified. If you do not remove the name DFSMRCL0 from IEAVTRML before you delete module DFSMRCL0 from SYS1.LPALIB, your z/OS system will not start.

If you previously used the AMASPZAP utility to zap DFSMRCL0 into the IEAVTRML CSECT (as is done in the IMS IVP), you must use the AMASPZAP utility to remove the name DFSMRCL0 from IEAVTRML. IEAVTRML is a table of 12-byte entries. The first 8 bytes of each entry is the name of the Resource Cleanup routine; the last 4 bytes must be zero. The last entry in the table must be all zeros, to indicate the end of the table. If DFSMRCL0 is not the last entry in the table, then in addition to removing the DFSMRCL0 entry, you must move any subsequent entries to ensure that no all-zero entries exist before the end of the table.

The following example shows how to remove DFSMRCL0 from IEAVTRML.

1. Use the AMASPZAP utility to dump the current contents of IEAVTRML:

```
//DMPVTRML JOB ...
//STEP001 EXEC PGM=AMASPZAP
//SYSLIB DD DSN=SYS1.LPALIB,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DUMP IGC0001C IEAVTRML
/*
```

2. Examine the contents of IEAVTRML from the AMASPZAP dump job output. Locate the entry containing DFSMRCL0 (in hex: X'C4C6E2D4D9C3D3F0'):

| **CCHHR- | 03C3000517 | RECORD   | LENGTH-  | 000F38   | MEMBER   | NAME     | IGC0001C | CSECT    | NAME | IEAVTRML |
|----------|------------|----------|----------|----------|----------|----------|----------|----------|------|----------|
| 000000   | C4C6E2D4   | D9C3D3F0 | 00000000 | D4E5D7E3 | E3D9D4D3 | 00000000 | C3C1D9F2 | D9E2C3F0 |      |          |
| 000020   | 00000000   | 00000000 | 00000000 | 00000000 | 00000000 | 00000000 | 00000000 | 00000000 |      |          |
| 000040   | 00000000   | 00000000 |          |          |          |          |          |          |      |          |

3. Use the AMASPZAP utility to replace the entry containing DFSMRCL0 with zeros. In the example output above, DFSMRCL0 is the first entry in IEAVTRML, and there are two other entries following it. To remove DFSMRCL0, entries 2 and 3 must be moved to become entries 1 and 2, and entry 3 must be zapped to be all zeros, as shown:

```
//DMPVTRML JOB ...
//STEP001 EXEC PGM=AMASPZAP
//SYSLIB DD DSN=SYS1.LPALIB,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
NAME IGC0001C IEAVTRML
VER 0000 C4C6E2D4D9C3D3F0
VER 000C D4E5D7E3E3D9D4D3
VER 0018 C3C1D9F2D9E2C3F0
REP 0000 D4E5D7E3E3D9D4D3
REP 000C C3C1D9F2D9E2C3F0
REP 0018 0000000000000000
/*
```

**Related reading:** For additional information, refer to *z/OS MVS Programming: Authorized Assembler Services Guide*.

## IMS abend formatting module

IMS installs the IMS abend formatting module (DFSAFMX0). No user setup is required, and you do not need to install module DFSAFMD0 on the host z/OS system.

Although module DFSAFMD0 is not required for IMS Version 11, this module is still shipped to support users who point to DFSAFMD0 directly in the IMS library. Module DFSAFMD0 that is shipped with IMS Version 11 can be used to provide abend formatting for IMS Version 10 and earlier.

If you want to use IMS online dump format support, you must have module DFSAFMD0 installed on any z/OS logical partition (LPAR) where you will run either online or batch IMS Version 10 or earlier. The version of module DFSAFMD0 that you install must be the same as the highest pre-Version 11 IMS that you will be running on the z/OS LPAR.

Module DFSAFMD0 processing is only for online dump formatting, which is done at the time of the abend. It has no effect on offline dump formatting, which is performed using an SDUMP or SYSMDUMP data set that is written at the time of the abend.

**Recommendation:** If you want to use online dump formatting, do not uninstall module DFSAFMD0 from the z/OS system until your migration to IMS Version 11 is complete, and there is no possibility that you will run an earlier version of IMS.

Module DFSAFMD0 and the DFSAFMX0 dynamic abend exit module installed by IMS Version 11 can coexist on the same system.

**Related reading:** For additional information about installing z/OS dump formatting exits, refer to *z/OS MVS Installation Exits*.

### Binding the abend formatting module (DFSAFMD0)

If you plan to run multiple versions of IMS at the same time, such as instances of IMS Version 9, IMS Version 10, and IMS Version 11 on the same logical partition (LPAR), and you want to use online dump formatting for IMS Version 9 and IMS Version 10, you must install the IMS abend formatting module (DFSAFMD0) as a z/OS exit in the IEAVADFM dump facility installation exit routine name list.

In IMS Version 10 and earlier, you must install module DFSAFMD0 on the host z/OS system as an IEAVADFM “Format SNAP, SYSABEND, and SYSUDUMP Dumps” exit.

If you plan to run IMS code from only IMS Version 11 and later on an LPAR, you do not need to install module DFSAFMD0.

To bind the module DFSAFMD0:

1. Bind module DFSAFMD0 into SYS1.LPALIB or an MLPA library.
2. Add the module name DFSAFMD0 to the IEAVADFM CSECT of the z/OS load module IGC0805A in SYS1.LPALIB. This process allows IMS module DFSAFMD0 to receive control during z/OS SNAP, SYSABEND, and SYSUDUMP processing so that it can format IMS control blocks in the dump.

**Related reading:** For additional information on the IEAVADFM installation process, refer to *z/OS MVS Installation Exits*.

## Uninstalling the abend formatting module (DFSAFMD0)

After you migrate to IMS Version 11 or later, you can remove module DFSAFMD0 from the host z/OS system.

To uninstall the module DFSAFMD0:

1. Remove the name DFSAFMD0 from the IEAVADFM CSECT of module IGC0805A in SYS1.LPALIB. Removing this name prevents the operating system from installing module DFSAFMD0 as an abend formatting exit at the next IPL.
2. Remove module DFSAFMD0 from SYS1.LPALIB or the MLPA library where module DFSAFMD0 was bound.
3. Restart with CLPA to enable these changes.

If you previously used the AMASPZAP utility to zap module DFSAFMD0 into the IEAVADFM CSECT (as is done in the IMS IVP), you must use the AMASPZAP utility to remove the name DFSAFMD0 from the IEAVADFM CSECT. The IEAVADFM CSECT is a table of 8-byte entries, followed by a final 4-byte entry that contains zeros to indicate the end of the exit name list. Each 8-byte entry contains the name of a dump formatting exit routine. If DFSAFMD0 is not the last entry in the table, then in addition to removing the DFSAFMD0 entry, you must move any subsequent entries to ensure that no all-zero entries exist before the end of the table.

The following example shows how to remove module DFSAFMD0 from the IEAVADFM CSECT.

1. Use the AMASPZAP utility to dump the current contents of the IEAVADFM CSECT:  

```
//DMPVADFM JOB ...
//STEP001 EXEC PGM=AMASPZAP
//SYSLIB DD DSN=SYS1.LPALIB,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DUMP IGC0805A IEAVADFM
/*
```
2. Examine the contents of the IEAVADFM CSECT from the AMASPZAP dump job output. Locate the entry that contains module DFSAFMD0 (in hexadecimal: X'C4C6E2C1C6D4C4F0'):

| RECORD | LENGTH   | MEMBER NAME | IGC0805A | CSECT NAME | IEAVADFM                            |
|--------|----------|-------------|----------|------------|-------------------------------------|
| 000000 | C4C6E2C1 | C6D4C4F0    | D4E8C4D4 | D7E7F0F0   | 00000000 00000000 00000000 00000000 |
| 000020 | 00000000 | 07FE0000    | 00000008 | 00000000   |                                     |

3. Use the AMASPZAP utility to replace the entry that contains module DFSAFMD0 with zeros. In the example output above, module DFSAFMD0 is the first entry in the IEAVADFM CSECT, and one other entry follows. To remove module DFSAFMD0, entry 2 must be moved to become entry 1, and entry 2 must be zapped to be all zeros, as shown:

```
/ZAPVADFM JOB ...
/STEP001 EXEC PGM=AMASPZAP
//SYSLIB DD DSN=SYS1.LPALIB,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
NAME IGC0805A IEAVADFM
VER 0000 C4C6E2C1C6D4C4F0
VER 0008 D4E8C4D4D7E7F0F0
REP 0000 D4E8C4D4D7E7F0F0
REP 0008 0000000000000000
```

## Adding the offline dump formatting routine to the print dump exit control table

Add the offline dump formatting module name to the Print Dump Exit Control Table in SYS1.PARMLIB member BLSCECT.

The entry must contain:

Module name DFSOFMD0

Exit flag 0

User verb IMSDUMP

An IMS Dump Formatter is also available from the component analysis section of the IPCS dialogs (IPCS ISPF selection 2.6).

If SDFSRESL is not in LNKSTxx, IPCS users must have SDFSRESL available in the JOBLIB or STEPLIB concatenation in order to be able to load DFSOFMD0.

### Related reading:

- For a description of the exit control table, see *z/OS MVS Initialization and Tuning Reference*.
- For more information about controlling IMS dumping options and about the Offline Dump Formatter, see *IMS Version 11 Diagnosis*.

## Binding the DBRC type 4 SVC

Bind the DBRC Type 4 SVC into an LPALIB or an MLPA library. It is named IGC00nnn, where nnn is the signed decimal SVC number.

## Authorizing IMS system data sets in the Authorized Program Facility

If you use JOBLIB/STEPLIB with region types of CTL (DB/DC region type), DBC (DBCTL region type), or DCC (DCCTL region type), all concatenations of the JOBLIB/STEPLIB must be APF authorized.

The following IMS system data sets must be APF authorized:

- IMS.SDXRRESL
- IMS.SDFSRESL
- IMS.SDFSJLIB
- IMS.MODBLKSA, IMS.MODBLKSB
- DFSESL DD, or the JOBLIB or STEPLIB into which your DB2 modules and tables are loaded

In addition to these data sets, in a DB/DC or DCCTL environment, SYS1.CSSLIB must be APF authorized. This is true regardless of whether you use APPC/z/OS. Even though SYS1.CSSLIB is in LNKSTxx and LNKSTxx is authorized, you must also have SYS1.CSSLIB in IEAAPFxx, because IMS accesses SYS1.CSSLIB without using the LNKSTxx concatenation. SYS1.CSSLIB must be explicitly APF-authorized.

**Recommendation:** Do not have the IMS.SDFSRESL in LNKSTxx when running multiple levels of IMS or when migrating to a new version or release level.



IMS conforms to z/OS rules for data set authorization. If you authorize an IMS job step, authorize all libraries used in that job step. To run an IMS batch region as non-authorized, concatenate a non-authorized library to IMS.SDFSRESL. To make this concatenation, the batch job must contain a DFSRESLB DD statement pointing to IMS.SDFSRESL.

**Related reading:** Refer to information on IEAAPFxx in *z/OS MVS Initialization and Tuning Reference*.

## Updating the APPC/MVS administration dialog

To use the APPC/MVS Administration Dialog utility with IMS TP Profiles, you must first add “IMS” as a transaction scheduler. To do this, you must add one line to the non-display panel ICQASE00 where the variable QASTSPE is defined.

The format of the line is as follows:

```
IMS,DFSTPPE0'
```

You must also change the single quotation mark (') on the current last line of the assignment to a plus sign (+).

In addition, IMS.SDFSEEXEC must be added to the TSO SYSPROC concatenation, and IMS.SDFSPLIB must be added to the TSO ISPLIB concatenation.

For more information on modifying this panel, see “Customizing the Dialog” in *z/OS MVS Planning: APPC/MVS Management*.

## RRS Archive Log Stream considerations

When using the z/OS Resource Recovery Services feature, be aware that the optional RRS Archive Log Stream has been found to generate large amounts of additional logging to the MVS logger.

If used, RRS Archive Logging should be monitored closely for system logstream performance impact. Neither IMS or RRS use the RRS Archive Log Stream data for diagnostic purposes, so it is not necessary to use this function.

## System Management Facility DDCONS parameter considerations

The setting of the System Management Facility (SMF) DDCONS parameter should be reviewed to avoid possible processing delays, as described in APAR OY31613.

The option is DDCONS. The two possible suboptions are YES and NO.

**YES** (Default) Requests that the consolidation be done

**NO** Requests that the consolidation not be done

---

## z/OS interface considerations for IRLM

There are many requirements that you must consider and required steps to ensure a complete and correct installation of IRLM on z/OS.

The following steps describe these requirements and required actions.



- Add the IRLM CTRACE module to the z/OS link list The IRLM CTRACE start/stop routine load module, DXRRL183, must reside in the z/OS Link List (LL). This module also contains the automatic restart manager (ARM) support for IRLM.
- Authorize IRLM in the Authorized Program Facility (APF) The IMS.SDXRRESL system data set must be APF authorized.
- Create IRLM subsystem names Unless you have deleted them, z/OS preconditioning has already defined IRLM and JRLM as subsystems names. You can use these names, or you can define your own. Create a z/OS subsystem name entry for each IRLM to be executed on the z/OS system. When two IRLMs reside in the same z/OS system, each must have a unique z/OS subsystem name.

Each message that the IRLM issues includes the IRLM z/OS subsystem name (IRLMNM on the start procedure) concatenated with the ID (IRLMID on the start procedure). A naming convention that allows easy identification of which IRLM issued a specific message is recommended. The following IRLM command displays all of the IRLM names and IDs associated with this IRLM or sharing group.

```
F irlmproc,STATUS,ALLI
```

- Add an IRLM entry to the z/OS Program Property Table (PPT) Unless you have deleted it, z/OS preconditioning has already defined a PPT entry for DXRRLM00.
- Update the print dump exit control table The entry must contain:
  - Module name DXRRLM50
  - Exit flag 0
  - User verb IRLM
 Ensure that one of these is true:
  - The print dump formatting module DXRRLM50 is in SYS1.LINKLIB.
  - The job that prints the dump contains a JOBLIB or STEPLIB statement specifying the library containing the modules.

**Related reading:**

- For more information about the Offline Dump Formatter, see *IMS Version 11 Diagnosis*.
- For information about adding an IRLM entry to the PPT, see “Adding the IRLM entry to the z/OS Program Properties Table” on page 129.
- Refer to z/OS MVS Initialization and Tuning Reference for information about:
  - Defining a subsystem to z/OS.
  - IEAAPFxx.
  - Responding to the messages and setting up PARMLIB members to contain trace options and parameters.

---

## Chapter 6. VTAM interface considerations

If your IMS system requires VTAM, the VTAM mode table must contain entries for all VTAM terminals defined to IMS.

You can use the table entry name at logon as any of the following:

- LOGMODE parameter on the VTAM VARY command
- MODE parameter on the /OPNDST command
- Parameter on the other terminal's INIT SELF command
- MODETBL parameter of the TERMINAL macro

The MODETBL parameter overrides any other entry supplied with the ACF/VTAM LOGON or SCIP exit CINIT. The MODETBL name for all parallel sessions with a given terminal is the same. Do not specify MODETBL for cross-domain resources.

The mode table entry creates the session parameters and thus controls the session established between IMS and the terminal.

**Related reading:** For a list of the BIND parameters for VTAM logical units, refer to *IMS Version 11 Communications and Connections*.

Define all of the following terminals:

- 3600, 3614, and SLU P as LUTYPE=0
- SLU 1 as LUTYPE=1
- SLU 2 as LUTYPE=2
- LU 6 as LUTYPE=6

A 3770P or 3790 terminal defined as SLUTYPE1 must be defined as unattended in its mode table entry. You can define a SLU 1 terminal as an exception or definite response for the secondary terminal. For terminals defined as SLUTYPEP, no options are allowed in the first 7 bytes of the BIND command.

When you specify PARSESS=NO in the VTAM APPL macro for IMS, VTAM parallel session support is not included. In this case, IMS counts as '1' within the MAXAPPL keyword of the VTAM START parameter.

When you specify PARSESS=YES in the VTAM APPL macro for IMS, VTAM parallel session support is included in the system. IMS counts as '2' within the MAXAPPL keyword of the VTAM START parameters.

**Related reading:** For information on IMS support for parallel sessions, see *IMS Version 11 Communications and Connections*. For more information on VTAM and Remote Site Recovery, see "RSR overview" on page 663.

**Important:** After the z/OS and VTAM interface steps are completed, you must start z/OS and specify either CLPA or MLPA=xx, or both.

---

## Setting the Network Control Program (NCP) delay

Set the value of the DELAY parameter on the HOST macro to 0 or as low as possible considering the other work in your system.

---

## Chapter 7. CSL administration

The tasks associated with administering a CSL are typically performed by the system administrator or system operator.

---

### Using the z/OS Automatic Restart Manager with the CSL

A CSL address space (ODBM, OM, RM, SCI), if requested, can register with the z/OS Automatic Restart Manager (ARM). The ARM is a z/OS recovery function that can improve the availability of started tasks. When a task fails or the system on which it is running fails, the ARM can restart the task without operator intervention.

IBM provides policy defaults for automatic restart management. You can use these defaults, or you can define your own ARM policy to specify how CSL address spaces should be restarted. The ARM policy specifies what should be done if the system fails or if a CSL address space fails.

To enable the ARM, you can specify ARMRST=Y in one of two ways:

- In the CSL address space initialization PROCLIB member data set
  - CSLDIxxx for ODBM
  - CSLOIxxx for OM
  - CSLRIxxx for RM
  - CSLSIxxx for SCI
- As an execution parameter

When ARM is enabled, the CSL address spaces register to ARM with an ARM element name, which are defined in the following table.

*Table 7. ARM element names*

| CSL address space | ARM element name        |
|-------------------|-------------------------|
| ODBM              | "CSL" + odbmname + "OD" |
| OM                | "CSL" + omname + "OM"   |
| RM                | "CSL" + rmname + "RM"   |
| SCI               | "CSL" + sciname + "SCI" |

**Note:** The name of the CSL address space is the name defined either as an execution parameter, or in the initialization PROCLIB member data set of that CSL address space. For example, if OMNAME=OM1A in the CSLOIxxx PROCLIB member data set, the ARM element name is CSLOM1AOM.

Use the appropriate ARM element name in your ARM policy for each CSL address space. For more information, see *z/OS MVS Setting Up a Sysplex*.

An abend table exists in the module for each CSL address space:

- CSLDARM0 for ODBM
- CSLOARM0 for OM
- CSLRARM0 for RM

- CSLSARM0 for SCI

The table lists the abends for which the ARM does not restart the CSL address space after the abend occurs. You can modify this table.

---

## Administering global online change

This topic covers enabling and disabling global online change for an IMSplex. This topic also covers maintaining an IMSplex with mixed online change scope.

### Enabling global online change

Enabling global online change does not require that you shut down all the IMS systems in the IMSplex at the same time. You can enable global online change one IMS at a time. After all of the IMS systems have switched to supporting global online change, the IMSplex only supports global online change.

**Restriction:** Global online change is not supported for RSR Tracker. If global online change is specified, the RSR tracking IMS abends.

Global online change requires a CSL. The following steps assume that the CSL environment is already established. To enable global online change, you must perform the following steps once:

1. Define the OLCSTAT data set characteristics.
2. Initialize OLCSTAT.

After you have defined and initialized the OLCSTAT data set, you can perform the following on one IMS at a time:

1. Remove MODSTAT DD statements from IMS control region JCL.
2. Remove MODSTAT2 DD statements from IMS control region JCL (if XRF).
3. Define DFSCGxxx PROCLIB member data set parameters related to global online change.
  - a. Specify OLC=GLOBAL to enable global online change.
  - b. Specify OLCSTAT= with the name of the online change data set. All IMS systems in the IMSplex must specify the same OLCSTAT data set.
4. Shut down IMS.
5. Cold start IMS.

When migrating batch IMS systems, you can enable Coordinated (global) Online Change by modifying the DBBATCH JCL:

- Add the OLCSTAT DD statement.
- Remove the MODSTAT DD statement.

### Disabling global online change

When disabling global online change you do not have to shut down all the IMS systems in the IMSplex at the same time. Also, after you disable global online change, the IMSplex only supports local online change.

When you are ready to disable global online change, perform the following steps on one IMS at a time:

1. Shut down IMS.
2. Modify the DFSCGxxx PROCLIB member data set parameters for online change to enable local online change by changing OLC=GLOBAL to OLC=LOCAL.

3. Run INITMOD job to initialize the MODSTAT data set.
4. Include MODSTAT DD statements to IMS control region JCL.
5. Include MODSTAT2 DD statements to IMS control region JCL (if XRF).
6. Cold start the IMS.

Typically, an IMS that was down during global online change has to cold start when it comes back up. Cold start might be required to prevent restart from processing log records against the current online change libraries that do not match, potentially causing restart to fail or causing a severe error later on. However, if the IMS was down for only the last global online change and the restart does not conflict with the last global online change, then the IMS is permitted to warm start.

If an IMS was down during only the last online change and its restart type conflicts with the last online change type, or IMS was down for two or more global online changes, IMS must cold start.

## Maintaining an IMSplex with mixed online change scope

If the IMSplex contains a mix of IMS systems that support global online change and those that do not, you have to manually coordinate online change with the global online change.

The IMS systems that are defined with OLC=GLOBAL can participate in the global online change.

For example, you can have a mixed IMSplex when enabling global online change across an IMSplex one IMS at a time. If you choose to perform online change during this migration, you have to manually coordinate online change on the IMS that specifies OLC=LOCAL with the global online change.

**Related reading:** For more information about the global online change process, see “Overview of the global online change function” on page 444.

---

## Administering automatic RECON loss notification

RECON loss notification is automatic after the IMSplex name for the RECON is specified.

Automatic RECON loss notification is activated as soon as one DBRC instance processes a RECON error. That error is immediately communicated to other DBRC instances through SCI. The DBRC instances that are notified of the RECON error issue a message (DSP1141I) indicating that they have been notified.

All DBRC instances that share a set of RECONs must also use the same IMSplex to ensure notification of RECON errors to all DBRC instances. The first use of RECON loss notification for a given set of RECONs sets the IMSplex name in the RECONs. Any subsequent attempts at RECON access using a different IMSplex or no IMSplex are rejected, and message DSP1136A is issued.

If you have ARLN active or parallel RECON access enabled, and you want to change the IMSplex associated with a set of RECONs, use the parameter IMSPLEX() | NOPLEX, which is on the CHANGE.RECON command. If any DBRC instances are active when the CHANGE.RECON IMSPLEX() | NOPLEX command

is executed, the command is rejected, and message DSP1137I is issued. The following list briefly summarizes the procedures to alter the IMSplex associated with a set of RECONS:

1. Wait for all DBRC activity on the current IMSplex to cease.
2. Submit a DBRC utility job to change the IMSplex name.
3. Alter the IMSplex name in the user exit routine, DSPSCIX0.
4. Ensure that the new IMSplex SCI is ready.
5. Resume DBRC activity on the new IMSplex.

**Note:** The IMSplex name in the RECON should only be changed when your SCI registration exit routine is changed to return the new IMSplex name.

To stop or reset automatic RECON loss notification, you must clear the IMSplex name. The following briefly summarizes the procedures to stop automatic RECON loss notification associated with a set of RECONS:

1. Wait for all DBRC activity on the current IMSplex to end.
2. Submit a DBRC utility job to change the IMSplex name.
  - a. Issue the command `CHANGE.RECON IMSPLEX() | NOPLEX`.

**Requirement:** If parallel RECON access is enabled and you intend to disable ARLN, you must disable parallel RECON access by issuing the `CHANGE.RECON ACCESS(SERIAL)` command prior to issuing the `CHANGE.RECON NOPLEX` command.

#### Related reading:

- For detailed information about automatic RECON loss notification, see “RECON loss notification” on page 562.
- For more information about parallel RECON access, see “Accessing the RECON data sets in parallel mode” on page 536.

---

## Monitoring an IMSplex

Monitoring the individual IMS systems in an IMSplex with CSL is not different from monitoring a standalone IMS system in any other environment. Monitoring the IMSplex as a whole, however, is different than monitoring an individual IMS system.

You can check system status by issuing the `QRY IMSPLEX SHOW(ALL)` command, which displays what IMS systems are active in the IMSplex and their statuses.

There are several other `QUERY` commands that are useful for monitoring certain aspects of the IMSplex. `QUERY` commands operate on an IMSplex-wide basis unless they are routed to specific IMS systems. For example:

- `QRY TRAN QNT(GT,nn)` where *nn* is a number you specify.  
The response displays the transactions that have a high queue count.
- `QRY DB STATUS(ALLOCFS,STOSCHD)`  
The response displays the databases that have an allocation failure or that have been stopped. There are various other statuses that can be specified.

---

## Diagnosing IMSplex problems

Each of the Common Service Layer address spaces, OM, RM, and SCI, produces SDUMPs for internal errors. You can find the CSL dumps in the SYS1.DUMP data sets.

To successfully diagnose IMSplex problems, you might need to collect one or more of the following types of system information:

- The SYSLOG from every logical partition (LPAR) where a CSL member resides.
- Output from the type-2 QUERY command to display the members of the IMSplex and their status.
- A z/OS SVC dump of the CSL address spaces.

You can write OM log records to a z/OS Logger data stream by specifying the AUDITLOG= parameter of the IMSPLEX keyword on the CSLOIxxx PROCLIB member data set. The log records provide an audit trail of command input, associated command output, and unsolicited message output.

You can view the messages in the audit trail using the TSO SPOC. Use the action bar to select **SPOC > Audit Trail** and fill in the data stream name and the time of interest.

For more information on diagnosing IMSplex problems, see *IMS Version 11 Diagnosis*.





---

## Chapter 8. CSL ODBM administration

The CSL Open Database Manager (ODBM) supports access to databases that are managed by IMS DB in DBCTL and DB/TM IMS systems in an IMSplex.

ODBM manages database connection and access requests from application programs that use the following resource adapters and APIs:

- IMS Universal Database resource adapter
- IMS Universal JDBC driver
- IMS Universal DL/I driver
- Open Database Access interface (ODBA)
- ODBM CSLDMI interface

Except for ODBA application programs and application programs written directly to the CSLDMI interface, all database connection and access requests are routed to ODBM from the resource adapters and APIs through IMS Connect. ODBM routes the database connection and access requests received from IMS Connect to the IMS systems that are managing the requested database.

The IMS Universal drivers that communicate with ODBM through IMS Connect use the open standard Distributed Relational Database Architecture (DRDA) as the low-level communication protocol. ODBM translates the incoming database requests from the Distributed Data Management (DDM) architecture that is a part of DRDA into the DL/I calls expected by IMS. When ODBM returns the IMS output to the client, ODBM translates the response back into the DDM protocol.

ODBA application programs and application programs that use the ODBM CSLDMI interface do not use DRDA but instead use the DL/I calls supported by the ODBA interface.

Existing ODBA applications can also use ODBM to protect IMS from abends caused by the unexpected termination of the ODBA applications during DL/I processing. ODBM support for ODBA application programs requires no changes to the application program itself and only minor changes to the DFSRP macro that defines the connections to IMS DB.

From the ODBM perspective, application programs that interface directly with ODBM through the CSLDMI interface, such as IMS Connect, are ODBM clients. Users can develop their own ODBM clients by using the ODBM CSLDMI interface. ODBM client application programs can access databases that are managed by IMS DB on any LPAR in an IMSplex

### **Related concepts:**

“Overview of the CSL Open Database Manager” on page 41

---

## ODBM client registration in a CSL

Before application programs can access IMS databases through ODBM, they, or their application server, must register with ODBM as a client. An ODBM client registers with ODBM by issuing the CSLDMREG request of the ODBM CSLDMI interface.

After the ODBM client has registered with ODBM, the application programs running under them can issue ODBA calls to IMS through ODBM.

Prior to registering with ODBM, the ODBM client must first register with SCI. When the ODBM client is done issuing ODBA calls to IMS, it can deregister by issuing a CSLDMDRG request.

For more information about the ODBM CSLDMI interface, including the CSLDMREG and CSLDMDRG requests, see *IMS Version 11 System Programming APIs*.

---

## ODBM and RRS

The CSL Open Database Manager (ODBM) can run with or without z/OS Resource Recovery Services. By default, ODBM runs with RRS.

When ODBM runs with RRS, RRS is the sync point manager and coordinates all the resource managers, including those other than IMS, that are associated with the UOR. ODBM uses the two-phase sync point protocol when running with RRS.

When ODBM runs without RRS, ODBM functions similarly to a CCTL that connects to IMS DB through the database resource adapter (DRA) interface; however, unlike a CICS CCTL, ODBM does not function as a syncpoint manager. Any syncpoint management must be performed by the ODBM client.

For more information about the DRA interface, see “The database resource adapter (DRA)” in *IMS Version 11 System Programming APIs*.

RRS support for ODBM is controlled by the RRS keyword in the ODBM initialization PROCLIB member CSLDIxxx, which is documented in *IMS Version 11 System Definition*.

---

## ODBM and ODBA

The Open Database Manager (ODBM) supports and uses the Open Database Access (ODBA) interface.

ODBM uses the ODBA interface to communicate with IMS. Many of the parameters in the ODBM CSLDCxxx configuration PROCLIB member are the same as those used by ODBA in the DFSxxxx0 startup table that is built from the DFSPRP macro parameters.

Because ODBM uses ODBA, ODBA application servers, such as DB2 for z/OS or WebSphere Application Server for z/OS can be configured to connect to IMS through ODBM instead of through ODBA. Connecting through ODBM prevents a U0113 abend from occurring if a DB2 stored procedure or WebSphere Application Server application program terminates unexpectedly during DL/I processing.

Configuring an ODBA application server to use ODBM does not require any changes to existing application programs that run under the ODBA application server.

## Configuring ODBA application servers to use ODBM

You can configure an application server that uses the Open Database Access (ODBA) interface, such as DB2 for z/OS or WebSphere Application Server for z/OS, to connect to IMS through the Open Database Manager (ODBM) component of the IMS Common Service Layer (CSL).

### Prerequisites:

To configure an application server that uses ODBA to connect to IMS through ODBM, you must configure the ODBA environment, as described in *Accessing IMS databases through the ODBA interface (Communications and Connections)*.

ODBM also requires the CSL Structured Call Interface (SCI) and the CSL Operations Manager (OM).

**Restriction:** The ODBA application server must reside on the same z/OS LPAR as both ODBM and IMS.

To configure an ODBA application server to use ODBM:

1. On the DFSPRP macro that is used to build the database resource adapter (DRA) startup table DFSxxxx0 (where xxxx is the alias name of the IMS system that is specified on the APSB call), in addition to the ODBA connection attributes, specify the IMSPLEX parameter and, optionally, the ODBMNAME parameter.

When the ODBMNAME parameter is not specified, ODBA selects the ODBM instance that the CSL Structured Call Interface (SCI) uses to route database connection requests.

2. Optional: Configure the ODBA application server to obtain and manage a private context under which the ODBA calls are processed. If the ODBA application server does not obtain a private context, ODBA obtains and manages the private context on behalf of the ODBA application.

### Related concepts:

 [CSL definition and tailoring \(System Definition\)](#)

### Related reference:

 [DRA startup table \(System Programming APIs\)](#)

---

## ODBM message routing

When you define the IMS systems in an IMSplex to ODBM, you also assign one or more alias names to each IMS system. The alias names are what the ODBM clients and ODBM use to route database access requests to the IMS systems.

If an alias name is shared by multiple IMS systems in an IMSplex, ODBM uses a round-robin algorithm to distribute incoming requests that specify the shared alias name among the sharing IMS systems.

Similarly, when incoming requests leave the alias name blank, ODBM also uses a round-robin algorithm to distribute incoming requests among all IMS systems available to ODBM in the IMSplex.

---

## ODBM and security

The CSL Open Database Manager (ODBM) does not perform user authentication or authorization itself.

If you are using IMS Connect with ODBM, you can authenticate the user IDs on request messages before they reach ODBM by using IMS Connect security. In addition to being able to call RACF directly, IMS Connect provides the IMS DB Security user exit routine (HWSAUTH0) to facilitate the customization of the security checking for communications with IMS DB.

You can have IMS check the authority of a user to allocate a PSB or access IMS resources by using APSB and resource access control (RAS) security.

APSB security is enabled by specifying the ODBASE parameter. RAS security is specified by the ISIS parameter. Both the ODBASE and ISIS parameters can be specified on the IMS or DBC startup procedure or the DFSPBxxx PROCLIB member.

If a user-written ODBM client passes a security object for a security product such as RACF, ODBM invokes RACROUTE REQUEST=VERIFY to create an accessor environment element (ACEE) for the APSB thread. IMS can then use the ACEE during APSB or RAS authorization for allocating PSBs or access to other resources.

### **Related concepts:**

Chapter 18, “IMS security,” on page 307

 Establishing and defining security (Communications and Connections)

---

## Chapter 9. CSL OM administration

This topic describes the administrative tasks associated with OM.

---

### CSL OM command routing

In an IMSplex environment, commands issued to OM can behave differently than commands issued to a single IMS system. Commands that are issued to OM are, by default, routed to all IMS systems in the IMSplex that are active and registered to process that particular command.

If you want to route a command to one or more specific command processing clients in the IMSplex, use the ROUTE parameter on the CSLOMCMMD request or the CSLOMI API. When an OM command has a ROUTE parameter, IMS chooses the highest level IMS in the route list as the command master. For example, in an IMSplex configuration that includes an IMS Version 11 system with an IMS Version 11 CQS, and another IMS Version 10 system with an IMS Version 10, if an INIT OLC command (which is a ROUTE=ANY command) is issued, the IMS Version 11 system is selected as the command master.

You can also specify routing information with the TSO SPOC and IMS Control Center. For information about how to specify the ROUTE parameter, see *IMS Version 11 System Programming APIs*. Refer also to the online help provided with the TSO SPOC and IMS Control Center.

In an IMSplex, a command originates from a SPOC or automated operator program (AOP) and is sent to OM. It is then routed to IMS systems; each IMS system returns its command output. OM then consolidates the individual responses and sends consolidated output, encapsulated in XML tags, back to the client originating the request. The following figure shows this routing. OS1 has an SCI, a SPOC TSO/ISPF application, and an automation program. OS2 has an OM with an SCI, and IMS control region with an SCI. OS3 has two IMS control regions, each with an SCI. The command is routed from OS1's SCI to OM in OS2. The SCI that is associated with that OM routes the command to other IMS systems through those system's SCI.




Figure 26. Command routing in an IMSplex with CSL

When commands are issued to an OM, command responses are encapsulated in XML tags.

For information about IMS commands and their responses, see *IMS Version 11 Commands, Volume 1: IMS Commands A-M* and *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

#### Related concepts:

- ➡ Interpreting CSL OM XML output (System Programming APIs)
- ➡ CSL OM automated operator program clients (System Programming APIs)

#### Related tasks:

- ➡ Registering an OM command processing client (System Programming APIs)

#### Related reference:

- ➡ Writing a CSL OM client (System Programming APIs)

## Issuing commands through the OM API

You can enter type-1 commands from an IMS terminal, such as the MTO terminal or an end-user terminal. You can enter many type-1 commands and all type-2 commands through the SPOC and OM API. When you enter a command through the OM API, OM can route the command to all IMS systems that can process the command. OM consolidates the command responses from the IMS systems into one buffer as they are received.

If the command has local scope only and is entered from the IMS Master Terminal, IMS end user terminal, or the IMS automated operator program (AOP), then the command is processed only by the local IMS where the command was entered. However, if a resource structure is part of the IMSplex, some commands, such as /STOP USER and /STOP NODE, have a global status in RM. If you enter those

commands from a SPOC, they might be routed to all of the IMS systems in the IMSplex and have a global effect. The TSO SPOC allows you to route commands to individual IMS systems.

CSL provides the type-2 command format. The benefits of this format are listed below:

- Elimination of resource name and command keyword conflicts
- Better command syntax checking using the BPE parser
- Simplified set of command verbs
- Parallel command processing
- Filters and wild cards for resource name selection
- Filtering and displaying of only selected output fields for QUERY command output
- Only one IMS with access to global information

Commands entered using this command format can only be entered through a SPOC. As a result, the command format cannot be entered through the system console, master terminal, end user terminal, DL/I call, APPC client, or through the OTMA interface. The format is based on simplified BPE parse rules. These commands work with the single-system image that an IMSplex provides of IMS systems sharing databases and queues in the IMSplex.

The type-2 command format is as follows. The keywords and parameters are optional:

ActionVerb ResourceType Keyword(*parameter*) Keyword(*parameter*) . . .

---

## CSL OM audit trail

An OM audit trail reflects the input and output of commands that are processed through the OM API. You can log: a point of command origin, what a particular command processed, and the result of that processing.

OM can write log records to a z/OS System Logger log stream, creating an audit trail of command input, associated command output, and unsolicited message output. Before using the OM audit trail function, the log stream needs to be created. Definitions for z/OS System Logger and coupling facility structures are required. To write OM log records to a z/OS log stream, specify the AUDITLOG= parameter of the IMSPLEX() keyword on the CSLOIxxx PROCLIB member.

If an individual command or command response plus its associated control data is longer than 32,760 bytes, it is split across multiple log records. These log records include control data that log formatting utilities can use to reassemble the log records into a single command or command response.

### Related concepts:

-  Other sample applications verified by the IVP (Installation)
-  OM audit trail (Operations and Automation)

### Related tasks:

-  OM audit log record formats (System Administration)



## Configuring the logstream for the CSL OM audit trail

The IMS IVP defines a logstream with parameters set to apply generally to most users. Depending on the configuration and requirements of your installation, you might consider changing these parameters when you define an OM audit trail logstream.

Parameters set by the IVP to consider changing include:

### **LS\_SIZE(200)**

LS\_SIZE controls the size of the offload data sets. Your installation might benefit from larger data sets.

### **LOWOFFLOAD(20)**

LOWOFFLOAD affects the frequency of offloads. A lower LOWOFFLOAD percentage, such as 0, generally means that offloads occur less frequently and, unless the logstream is browsed in realtime, should be the most efficient.

### **HIGHOFFLOAD(50)**

HIGHOFFLOAD affects the frequency of offloads. A higher HIGHOFFLOAD value, such as 70, generally means that offloads occur less frequently, which improves efficiency.

The IMS IVP also does not define the system logger to duplex to staging data sets. If your installation requires a higher level of data redundancy than is provided with MVS dataspace duplexing then consider enabling external duplexing by adding STG\_DUPLEX(YES) to the definition below, along with the DUPLEXMODE parameter, which defines the conditions in which the duplexing occurs. However, the higher level of data redundancy must be weighed against the performance impact of the duplexing option.

The IMS IVP defines a logstream with the following parameters:

```
DEFINE STRUCTURE NAME(IMSOM2Q01)
    LOGSNUM(1)
    AVGBUFSIZE(4000)
    MAXBUFSIZE(32760)

DEFINE LOGSTREAM NAME(SYSLOG.OM2Q01.LOG)
    STRUCTNAME(IMSOM2Q01)
    HLQ(IXGLOGR)
    LS_STORCLAS(LOGGER1)
    LS_DATACLAS(LOGGER1)
    LS_MGMTCLAS(LOGGER)
    LS_SIZE(200)
    LOWOFFLOAD(20)
    HIGHOFFLOAD(50)
```

### **Recommendations:**

- Set the logstream maximum buffer size (MAXBUFSIZE) to 32,760 bytes. This setting ensures that all command and command response data is written to the audit trail. A lower setting will result in the truncation of some OM Audit Trail log records.
- Dedicate a log stream to the OM audit trail.

### **Related concepts:**

 IBM Redbooks: System Logger

 z/OS: Using system logger services

### **Related tasks:**

➡ z/OS: Planning for system logger applications

**Related reference:**

➡ CSLOIxxx member of the IMS PROCLIB data set (System Definition)

## OM audit log record formats

Each Operations Manager (OM) audit log record contains a log record prefix, followed by data that is unique to the record.

Use macro CSLZLGPF to map the log record prefix. To view the OM audit log record formats, assemble mapping macro CSLOLGRC.

Individual log record DSECTs can be obtained by assembling the ILOGREG macro while including a format statement for the desired log record.

The following table shows the OM audit log records. For each OM audit log record, the table shows:

- The log record type and subtype
- The macro that maps the record
- The event that causes the record to be written

*Table 8. OM audit log records*

| Type  | Subtype | Mapping macro | Event that causes the log record to be written                                                                                                                |
|-------|---------|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| X'06' | X'01'   | CSLOLGCM      | The command input was logged before calling the OM Input user exit routine.                                                                                   |
| X'06' | X'02'   | CSLOLGCM      | The command input was logged after calling the OM Input user exit routine and after that user exit routine modified the command input.                        |
| X'08' | X'01'   | CSLOLGCR      | The command response output was logged before calling the OM Output user exit routine.                                                                        |
| X'08' | X'02'   | CSLOLGCR      | The command response output was logged after calling the OM Output user exit routine and after that user exit routine modified the command response output.   |
| X'09' | X'01'   | CSLOLGOU      | The unsolicited output message was logged before calling the OM Output user exit routine.                                                                     |
| X'09' | X'02'   | CSLOLGOU      | The unsolicited output message was logged after calling the OM Output user routine, and after that user exit routine modified the unsolicited output message. |

## Printing OM log records

To print the OM log records from the z/OS system log, use the IMS File Select and Formatting Print utility (DFSERA10) with exit routine CSLOERA3.

The example in the following figure shows the JCL that is required to print the log records from a z/OS system log.

### *JCL to print OM log records from a z/OS system log*

```
//CSLERA1 JOB MSGLEVEL=1,MSGCLASS=A,CLASS=K
//STEP1 EXEC PGM=DFSERA10
//STEPLIB DD DISP=SHR,DSN=IMS.SDFSRESL
//SYSPRINT DD SYSOUT=A
//SYSUT1 DD DSN=SYSLOG.OM.AUDIT.TRAIL.LOG,
// SUBSYS=(LOGR,IXGSEXIT),
// DCB=(BLKSIZE=32760)
```

```
//SYSIN      DD
CONTROL      CNTL      H=EOF
OPTION       PRINT     EXITR=CSLULALE
END
//
```

Command input and command response log records longer than 32 760 bytes are divided into two or more segments. Unsolicited output has a maximum length of 320 bytes. Each segment is sent to the MVS logger on a separate IXGWRITE call from OM. Therefore, long log records appear as multiple records on the log stream. Additional control data is added to the command response and command input log records to enable the CSLOERA3 and CSLULALE formatting utilities to reassemble log records that have been separated into multiple segments. A bit is added to the log record prefix to indicate single segment log records, including log records other than command input or command response. Log records are printed in the same order that they were sent to the log stream, based on the timestamp provided in the log record. If MAXBUFSIZE is set lower than 32 760 bytes, log records will not be broken into multiple segments. Instead, any log record that exceeds the length in MAXBUFSIZE will be truncated.

**Recommendation:** Set the MAXBUFSIZE to 32760 bytes. Log records requiring more than 65,535 segments will continue to be truncated. Currently, this limitation affects only a log record longer than 2GB.

The DD statements for printing OM log records are:

#### STEPLIB

DSN= points to IMS.SDFSRESL, which contains the IMS File Select and Formatting Print utility, DFSERA10.

#### SYSUT1

DSN= points to the OM log stream name that is specified on the AUDITLOG= parameter in the CSLOIxxx PROCLIB member.

The control statements for printing the OM log records are:

**H=** Specifies the number of log records to print. To print all of the log records, specify H=EOF.

#### EXITR=CSLULALE

Identifies the OM log record exit routine that is called to format each log record. Exit routine CSLULALE formats the OM audit trail in a format similar to a syslog. Exit routine CSLOERA3 can also be used. Exit routine CSLOERA3 prints the records in a dump format, including the record type and time-stamp information for each record, and provides the contents of the record up to a maximum of 32 760 bytes.

To limit the log data to a specific time range, use the FROM and TO parameters on the SUBSYS statement, as shown in the figure below. This DD card prints log records from 11:00 to 12:00 on day 42 of the year 2007.

#### *JCL to limit log data to a specific time range*

```
//SYSUT1      DD      DSN=SYSLOG.OM.AUDIT.TRAIL.LOG,
//              SUBSYS=(LOGR,IXGSEXIT,
//              'FROM=(2007/042,11:00:00),TO=(2007/042,12:00:00)'),
//              DCB=(BLKSIZE=32760)
```

Dates and times are specified in Greenwich Mean Time (GMT). The seconds field in the time value is optional. To use local dates and times, add the LOCAL keyword, as shown in the following table.

*JCL to limit log data to a specific time range, using local dates and times*

```
//SYSUT1 DD DSN=SYSLOG.OM.AUDIT.TRAIL.LOG,
//          SUBSYS=(LOGR,IXGSEXIT,
//          'FROM=(2007/042,11:00:00),TO=(2007/042,12:00:00),LOCAL'),
//          DCB=(BLKSIZE=32760)
```

**Related reference:**

 File Select and Formatting Print utility (DFSERA10) (System Utilities)

## CSL OM command security

OM command security is optionally performed during command processing.

Command security allows:

- The user to control which user IDs can enter IMS commands through OM
- The user ID to be associated with an application program address space
- The user ID to be the end user logged onto TSO SPOC

The CMDSEC= parameter is available on the OM startup procedure (CSLOM), the OM initialization PROCLIB member data set (CSLOIxxx), and the DFSCGxxx PROCLIB member data set. When it is issued as part of the OM startup procedure, it applies to all IMS commands, type-1 and type-2. When it is issued using the DFSCGxxx PROCLIB member data set, it applies only to type-1 commands entered through OM. The differences in OM and IMS security are described in the following table.

*Table 9. Comparing OM and IMS security*

| Security method                             | N                                                                                                    | A                                                                                                    | E                                                                                      | R                                                                                       |
|---------------------------------------------|------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| OM Execution Parameter (CSLOM and CSLOIxxx) | No authorization checking is performed. This is the default.                                         | Calls both RACF and the CSL OM Security user exit routine for command authorization.                 | Calls the CSL OM Security user exit routine for command authorization.                 | Calls RACF for command authorization. Commands are part of the OPERCMDS resource class. |
| DFSCGxxx PROCLIB member data set            | No authorization checking is performed. This is the default. OM might perform command authorization. | Calls both RACF and the IMS Command Authorization Exit routine (DFSCCMD0) for command authorization. | Calls the IMS Command Authorization Exit routine (DFSCCMD0) for command authorization. | Calls RACF for command authorization. Commands are part of the CIMS resource class.     |

**Recommendation:** Use OM command security rather than IMS command security.

RACF access authorities (READ or UPDATE) and resource names for all commands supported through the OM API are described in “IMS commands, RACF access authorities and resource names table” on page 159. The RACF authorities indicate the access authority with which the command was registered.

Commands are registered to OM with the CSLOMBLD request. The access authority on the RACF PERMIT command must match the access authority with which the command was registered. For more information on registering commands with CSLOMBLD, see *IMS Version 11 System Programming APIs*.

## RACF authorization for commands entered through OM

All OM security checking uses the RACF OPERCMDS class. The resource name is in the form of: *IMS.imsplexname.commandverb.commandkeyword*

### IMS

The high-level qualifier for all RACF resources that protect IMS commands entered through OM.

### *imsplexname*

The name of the IMSplex for which the command is authorized. The IMSplex name must include the characters CSL at the beginning. For example, CSLplex1.

### *commandverb*

The name of the command verb.

### *commandkeyword*

The primary command keyword or resource type for the command.

## Example: authorizing users to commands in an IMSplex

The following example shows a RACF definition that authorizes various users to different commands.

```
/* allow UPD TRAN command on any IMSplex for user ID Jim */
RDEFINE OPERCMDS IMS.*.UPD.TRAN UACC(NONE)
PERMIT IMS.*.UPD.TRAN CLASS(OPERCMDS) ID(JIM) ACCESS(UPDATE)

/* allow INIT OLC command on IMSplex CSLPLEXA for user ID Sandy */
RDEFINE OPERCMDS IMS.CSLPLEXA.INIT.OLC UACC(NONE)
PERMIT IMS.CSLPLEXA.INIT.OLC CLASS(OPERCMDS) ID(SANDY) ACCESS(UPDATE)

/* allow any QRY command on any IMSplex for user ID Tom */
RDEFINE OPERCMDS IMS.*.QRY.* UACC(NONE)
PERMIT IMS.*.QRY.* CLASS(OPERCMDS) ID(TOM) ACCESS(READ)

/* allow any commands on any IMSplex for user ID Betty */
RDEFINE OPERCMDS IMS.* UACC(NONE)
PERMIT IMS.* CLASS(OPERCMDS) ID(BETTY) ACCESS(UPDATE)

SETROPTS CLASSACT(OPERCMDS)
SETROPTS RACLIST(OPERCMDS) REFRESH
```

## Example: using RACF groups and the OPERCMDS class for type-2 command security

The following example shows a RACF definition that authorizes all user IDs that are part of a RACF group to use any IMS operator commands. This example shows two important elements of RACF administration:

- RACF groups
- Wildcard characters

```
/* define IMS operator group */
ADDGROUP IMSOPER
CONNECT (PEDRO,ROSA,PETER,MATT) GROUP(IMSOPER)

/* allow any commands on any IMSplex for group IMSOPER */
```

```
RDEFINE OPERCMDS IMS.* UACC(NONE)
PERMIT IMS.* CLASS(OPERCMDS) ID(IMSOPER) ACCESS(UPDATE)

SETOPTS CLASSACT(OPERCMDS)
SETOPTS RACLIST(OPERCMDS) REFRESH
```

Using a group definition simplifies the administration of RACF profiles. The group can be permitted to numerous RACF profiles, such as data sets and OPERCMDS profiles. If a user changes jobs, you need to remove the user ID only from the group definition instead of deleting its access to each RACF resource profile.

The benefit of using wildcard characters in the RACF resource name is that you can keep the number of RACF resource definitions to a minimum. Also, any matching operator commands that are defined after the OPERCMDS profiles are created are still protected from unspecified users.

### Example: authorizing users who are associated with the RACF APPL class

The following example shows a RACF definition that adds a profile for authorizing users who are associated with the RACF APPL class.

```
/* add a profile */
RDEFINE APPL JOBNAME UACC(NONE)
PERMIT JOBNAME CLASS(APPL) ID(USERID OR GROUPNAME) ACCESS(UPDATE)
```

## IMS commands, RACF access authorities and resource names table

The following table applies to the OM command security for IMS type-1 and type-2 command security. The table lists, by IMS command verb and keyword, the resource name and authorization that are used for RACF security checking.

For example, if you want to issue the ACTIVATE NODE command on an IMSplex named PLX01, your user ID must have UPDATE access to profile IMS.CSLPLX01.ACT.NODE in the OPERCMDS class. If you only have READ authority, the ACTIVATE NODE command will fail with an 'insufficient authority' message.

In general, you can issue a display or query a command with READ authority. To change the state of a resource, you need UPDATE authority.

**Note:** You can reduce the number of RACF resources required by using wildcards in the resource name. For example, use IMS.CSLPLX01.DIS.\* to grant authority to use all display commands.

Table 10. Resource names and RACF authority for IMS commands

| Command verb | Command keyword | Authority | Resource name         |
|--------------|-----------------|-----------|-----------------------|
| ACT          | LINK            | UPDATE    | IMS.plxname.ACT.LINK  |
| ACT          | NODE            | UPDATE    | IMS.plxname.ACT.NODE  |
| ALL          | LU              | UPDATE    | IMS.plxname.ALL.LU    |
| ASS          | CLASS           | UPDATE    | IMS.plxname.ASS.CLASS |
| ASS          | CPRI            | UPDATE    | IMS.plxname.ASS.CPRI  |
| ASS          | INPUT           | UPDATE    | IMS.plxname.ASS.INPUT |
| ASS          | LCT             | UPDATE    | IMS.plxname.ASS.LCT   |

Table 10. Resource names and RACF authority for IMS commands (continued)

| Command verb | Command keyword | Authority | Resource name              |
|--------------|-----------------|-----------|----------------------------|
| ASS          | LPRI            | UPDATE    | IMS.plxname.ASS.LPRI       |
| ASS          | LTERM           | UPDATE    | IMS.plxname.ASS.LTERM      |
| ASS          | NPRI            | UPDATE    | IMS.plxname.ASS.NPRI       |
| ASS          | OUTPUT          | UPDATE    | IMS.plxname.ASS.OUTPUT     |
| ASS          | PARLIM          | UPDATE    | IMS.plxname.ASS.PARLIM     |
| ASS          | PLCT            | UPDATE    | IMS.plxname.ASS.PLCT       |
| ASS          | SEGNO           | UPDATE    | IMS.plxname.ASS.SEGNO      |
| ASS          | SEGSZ           | UPDATE    | IMS.plxname.ASS.SEGSZ      |
| ASS          | TRAN            | UPDATE    | IMS.plxname.ASS.TRAN       |
| ASS          | USER            | UPDATE    | IMS.plxname.ASS.USER       |
| BRO          | ACT             | READ      | IMS.plxname.BRO.ACT        |
| BRO          | LINE            | READ      | IMS.plxname.BRO.LINE       |
| BRO          | LTERM           | READ      | IMS.plxname.BRO.LTERM      |
| BRO          | MASTER          | READ      | IMS.plxname.BRO.MASTER     |
| BRO          | NODE            | READ      | IMS.plxname.BRO.NODE       |
| BRO          | PTERM           | READ      | IMS.plxname.BRO.PTERM      |
| BRO          | USER            | READ      | IMS.plxname.BRO.USER       |
| CHA          | APPC            | UPDATE    | IMS.plxname.CHA.APPC       |
| CHA          | CCTL            | UPDATE    | IMS.plxname.CHA.CCTL       |
| CHA          | CPLOG           | UPDATE    | IMS.plxname.CHA.CPLOG      |
| CHA          | DESC            | UPDATE    | IMS.plxname.CHA.DESC       |
| CHA          | DIR             | UPDATE    | IMS.plxname.CHA.DIR        |
| CHA          | FDR             | UPDATE    | IMS.plxname.CHA.FDR        |
| CHA          | LINK            | UPDATE    | IMS.plxname.CHA.LINK       |
| CHA          | NODE            | UPDATE    | IMS.plxname.CHA.NODE       |
| CHA          | PSWD            | UPDATE    | IMS.plxname.CHA.PSWD       |
| CHA          | SUBSYS          | UPDATE    | IMS.plxname.CHA.SUBSYS     |
| CHA          | SURV            | UPDATE    | IMS.plxname.CHA.SURV       |
| CHA          | TRAN            | UPDATE    | IMS.plxname.CHA.TRAN       |
| CHA          | UOR             | UPDATE    | IMS.plxname.CHA.UOR        |
| CHA          | USER            | UPDATE    | IMS.plxname.CHA.USER       |
| CHE          |                 | UPDATE    | IMS.plxname.CHE            |
| CHE          | DUMPQ           | UPDATE    | IMS.plxname.CHE.DUMPQ      |
| CHE          | FREEZE          | UPDATE    | IMS.plxname.CHE.FREEZE     |
| CHE          | PURGE           | UPDATE    | IMS.plxname.CHE.PURGE      |
| CHE          | STATISTICS      | UPDATE    | IMS.plxname.CHE.STATISTICS |
| CLS          | NODE            | UPDATE    | IMS.plxname.CLS.NODE       |
| CQC          | SHRQ            | UPDATE    | IMS.plxname.CQC.SHRQ       |
| CQC          | SYSTEM          | UPDATE    | IMS.plxname.CQC.SYSTEM     |
| CQQ          | STATISTICS      | READ      | IMS.plxname.CQQ.STATISTICS |



Table 10. Resource names and RACF authority for IMS commands (continued)

| Command verb | Command keyword | Authority | Resource name            |
|--------------|-----------------|-----------|--------------------------|
| CQS          | SHUTDOWN        | UPDATE    | IMS.plxname.CQS.SHUTDOWN |
| CRE          | DB              | UPDATE    | IMS.plxname.CRE.DB       |
| CRE          | DBDESC          | UPDATE    | IMS.plxname.CRE.DBDESC   |
| CRE          | PGM             | UPDATE    | IMS.plxname.CRE.PGM      |
| CRE          | PGMDESC         | UPDATE    | IMS.plxname.CRE.PGMDESC  |
| CRE          | RTC             | UPDATE    | IMS.plxname.CRE.RTC      |
| CRE          | RTCDESC         | UPDATE    | IMS.plxname.CRE.RTCDESC  |
| CRE          | TRAN            | UPDATE    | IMS.plxname.CRE.TRAN     |
| CRE          | TRANDESC        | UPDATE    | IMS.plxname.CRE.TRANDESC |
| DBD          | DB              | UPDATE    | IMS.plxname.DBD.DB       |
| DBR          | AREA            | UPDATE    | IMS.plxname.DBR.AREA     |
| DBR          | DB              | UPDATE    | IMS.plxname.DBR.DB       |
| DBR          | DATAGRP         | UPDATE    | IMS.plxname.DBR.DATAGRP  |
| DEL          | DB              | UPDATE    | IMS.plxname.DEL.DB       |
| DEL          | DBDESC          | UPDATE    | IMS.plxname.DEL.DBDESC   |
| DEL          | LE              | UPDATE    | IMS.plxname.DEL.LE       |
| DEL          | PGM             | UPDATE    | IMS.plxname.DEL.PGM      |
| DEL          | PGMDESC         | UPDATE    | IMS.plxname.DEL.PGMDESC  |
| DEL          | PSWD            | UPDATE    | IMS.plxname.DEL.PSWD     |
| DEL          | RTC             | UPDATE    | IMS.plxname.DEL.RTC      |
| DEL          | RTCDESC         | UPDATE    | IMS.plxname.DEL.RTCDESC  |
| DEL          | TERMINAL        | UPDATE    | IMS.plxname.DEL.TERMINAL |
| DEL          | TRAN            | UPDATE    | IMS.plxname.DEL.TRAN     |
| DEL          | TRANDESC        | UPDATE    | IMS.plxname.DEL.TRANDESC |
| DEQ          | AOITKN          | UPDATE    | IMS.plxname.DEQ.AOITKN   |
| DEQ          | LINE            | UPDATE    | IMS.plxname.DEQ.LINE     |
| DEQ          | LTERM           | UPDATE    | IMS.plxname.DEQ.LTERM    |
| DEQ          | LU              | UPDATE    | IMS.plxname.DEQ.LU       |
| DEQ          | MSNAME          | UPDATE    | IMS.plxname.DEQ.MSNAME   |
| DEQ          | NODE            | UPDATE    | IMS.plxname.DEQ.NODE     |
| DEQ          | SUSPEND         | UPDATE    | IMS.plxname.DEQ.SUSPEND  |
| DEQ          | TMEM            | UPDATE    | IMS.plxname.DEQ.TMEM     |
| DEQ          | TRAN            | UPDATE    | IMS.plxname.DEQ.TRAN     |
| DEQ          | USER            | UPDATE    | IMS.plxname.DEQ.USER     |
| DIAG         | ADDRESS         | UPDATE    | IMS.plxname.DIAG.ADDRESS |
| DIAG         | BLOCK           | UPDATE    | IMS.plxname.DIAG.BLOCK   |
| DIAG         | LTERM           | UPDATE    | IMS.plxname.DIAG.LTERM   |
| DIAG         | NODE            | UPDATE    | IMS.plxname.DIAG.NODE    |
| DIAG         | SNAP            | UPDATE    | IMS.plxname.DIAG.SNAP    |
| DIAG         | TRAN            | UPDATE    | IMS.plxname.DIAG.TRAN    |



Table 10. Resource names and RACF authority for IMS commands (continued)

| Command verb | Command keyword | Authority | Resource name             |
|--------------|-----------------|-----------|---------------------------|
| DIAG         | USER            | UPDATE    | IMS.plxname.DIAG.USER     |
| DIS          | ACT             | READ      | IMS.plxname.DIS.ACT       |
| DIS          | AFFIN           | READ      | IMS.plxname.DIS.AFFIN     |
| DIS          | AOITKN          | READ      | IMS.plxname.DIS.AOITKN    |
| DIS          | APPC            | READ      | IMS.plxname.DIS.APPC      |
| DIS          | AREA            | READ      | IMS.plxname.DIS.AREA      |
| DIS          | ASMT            | READ      | IMS.plxname.DIS.ASMT      |
| DIS          | CCTL            | READ      | IMS.plxname.DIS.CCTL      |
| DIS          | CONV            | READ      | IMS.plxname.DIS.CONV      |
| DIS          | CPLOG           | READ      | IMS.plxname.DIS.CPLOG     |
| DIS          | CQS             | READ      | IMS.plxname.DIS.CQS       |
| DIS          | DB              | READ      | IMS.plxname.DIS.DB        |
| DIS          | DBD             | READ      | IMS.plxname.DIS.DBD       |
| DIS          | DESC            | READ      | IMS.plxname.DIS.DESC      |
| DIS          | FDR             | READ      | IMS.plxname.DIS.FDR       |
| DIS          | FPV             | READ      | IMS.plxname.DIS.FPV       |
| DIS          | HSB             | READ      | IMS.plxname.DIS.HSB       |
| DIS          | HSSP            | READ      | IMS.plxname.DIS.HSSP      |
| DIS          | LINE            | READ      | IMS.plxname.DIS.LINE      |
| DIS          | LINK            | READ      | IMS.plxname.DIS.LINK      |
| DIS          | LTERM           | READ      | IMS.plxname.DIS.LTERM     |
| DIS          | LU              | READ      | IMS.plxname.DIS.LU        |
| DIS          | MASTER          | READ      | IMS.plxname.DIS.MASTER    |
| DIS          | MODIFY          | READ      | IMS.plxname.DIS.MODIFY    |
| DIS          | MSNAME          | READ      | IMS.plxname.DIS.MSNAME    |
| DIS          | NODE            | READ      | IMS.plxname.DIS.NODE      |
| DIS          | OASN            | READ      | IMS.plxname.DIS.OASN      |
| DIS          | OLDS            | READ      | IMS.plxname.DIS.OLDS      |
| DIS          | OTMA            | READ      | IMS.plxname.DIS.OTMA      |
| DIS          | OVERFLOWQ       | READ      | IMS.plxname.DIS.OVERFLOWQ |
| DIS          | POOL            | READ      | IMS.plxname.DIS.POOL      |
| DIS          | PGM             | READ      | IMS.plxname.DIS.PGM       |
| DIS          | PSB             | READ      | IMS.plxname.DIS.PSB       |
| DIS          | PTERM           | READ      | IMS.plxname.DIS.PTERM     |
| DIS          | Q               | READ      | IMS.plxname.DIS.Q         |
| DIS          | QCNT            | READ      | IMS.plxname.DIS.QCNT      |
| DIS          | RECOVERY        | READ      | IMS.plxname.DIS.RECOVERY  |
| DIS          | RTC             | READ      | IMS.plxname.DIS.RTC       |
| DIS          | SHUTDOWN        | READ      | IMS.plxname.DIS.SHUTDOWN  |
| DIS          | STATUS          | READ      | IMS.plxname.DIS.STATUS    |

Table 10. Resource names and RACF authority for IMS commands (continued)

| Command verb | Command keyword | Authority | Resource name            |
|--------------|-----------------|-----------|--------------------------|
| DIS          | STRUC           | READ      | IMS.plxname.DIS.STRUC    |
| DIS          | SUBSYS          | READ      | IMS.plxname.DIS.SUBSYS   |
| DIS          | SYSID           | READ      | IMS.plxname.DIS.SYSID    |
| DIS          | TIMEOVER        | READ      | IMS.plxname.DIS.TIMEOVER |
| DIS          | TMEM            | READ      | IMS.plxname.DIS.TMEM     |
| DIS          | TRACE           | READ      | IMS.plxname.DIS.TRACE    |
| DIS          | TRACKING        | READ      | IMS.plxname.DIS.TRACKING |
| DIS          | TRAN            | READ      | IMS.plxname.DIS.TRAN     |
| DIS          | UOR             | READ      | IMS.plxname.DIS.UOR      |
| DIS          | USER            | READ      | IMS.plxname.DIS.USER     |
| END          | LINE            | UPDATE    | IMS.plxname.END.LINE     |
| END          | NODE            | UPDATE    | IMS.plxname.END.NODE     |
| END          | USER            | UPDATE    | IMS.plxname.END.USER     |
| ERE          |                 | UPDATE    | IMS.plxname.ERE          |
| ERE          | BACKUP          | UPDATE    | IMS.plxname.ERE.BACKUP   |
| ERE          | COLDBASE        | UPDATE    | IMS.plxname.ERE.COLDBASE |
| ERE          | COLDCOMM        | UPDATE    | IMS.plxname.ERE.COLDCOMM |
| ERE          | COLDSYS         | UPDATE    | IMS.plxname.ERE.COLDSYS  |
| EXC          | LINE            | UPDATE    | IMS.plxname.EXC.LINE     |
| EXC          | NODE            | UPDATE    | IMS.plxname.EXC.NODE     |
| EXC          | USER            | UPDATE    | IMS.plxname.EXC.USER     |
| EXI          | CONV            | UPDATE    | IMS.plxname.EXI.CONV     |
| EXP          | DEFN            | UPDATE    | IMS.plxname.EXP.DEFN     |
| IDL          | LINE            | UPDATE    | IMS.plxname.IDL.LINE     |
| IDL          | LINK            | UPDATE    | IMS.plxname.IDL.LINK     |
| IDL          | NODE            | UPDATE    | IMS.plxname.IDL.NODE     |
| IMP          | DEFN            | UPDATE    | IMS.plxname.IMP.DEFN     |
| INIT         | OLC             | UPDATE    | IMS.plxname.INIT.OLC     |
| INIT         | OLREORG         | UPDATE    | IMS.plxname.INIT.OLREORG |
| LOC          | DB              | UPDATE    | IMS.plxname.LOC.DB       |
| LOC          | PGM             | UPDATE    | IMS.plxname.LOC.PGM      |
| LOC          | TRAN            | UPDATE    | IMS.plxname.LOC.TRAN     |
| LOG          |                 | UPDATE    | IMS.plxname.LOG          |
| MOD          | ABORT           | UPDATE    | IMS.plxname.MOD.ABORT    |
| MOD          | COMMIT          | UPDATE    | IMS.plxname.MOD.COMMIT   |
| MOD          | PREPARE         | UPDATE    | IMS.plxname.MOD.PREPARE  |
| MON          | LINE            | UPDATE    | IMS.plxname.MON.LINE     |
| MSA          | LINK            | UPDATE    | IMS.plxname.MSA.LINK     |
| MSA          | MSNAME          | UPDATE    | IMS.plxname.MSA.MSNAME   |
| MSA          | SYSID           | UPDATE    | IMS.plxname.MSA.SYSID    |

Table 10. Resource names and RACF authority for IMS commands (continued)

| Command verb | Command keyword | Authority | Resource name             |
|--------------|-----------------|-----------|---------------------------|
| MSA          | TRAN            | UPDATE    | IMS.plxname.MSA.TRAN      |
| MSV          | MSNAME          | UPDATE    | IMS.plxname.MSV.MSNAME    |
| MSV          | SYSID           | UPDATE    | IMS.plxname.MSV.SYSID     |
| NRE          |                 | UPDATE    | IMS.plxname.NRE           |
| NRE          | CHKPT           | UPDATE    | IMS.plxname.NRE.CHKPT     |
| OPN          | NODE            | UPDATE    | IMS.plxname.OPN.NODE      |
| PST          | LINE            | UPDATE    | IMS.plxname.PST.LINE      |
| PST          | LINK            | UPDATE    | IMS.plxname.PST.LINK      |
| PST          | LTERM           | UPDATE    | IMS.plxname.PST.LTERM     |
| PST          | MSPLINK         | UPDATE    | IMS.plxname.PST.MSPLINK   |
| PST          | REGION          | UPDATE    | IMS.plxname.PST.REGION    |
| PST          | TRAN            | UPDATE    | IMS.plxname.PST.TRAN      |
| PUR          | APPC            | UPDATE    | IMS.plxname.PUR.APPC      |
| PUR          | FPPROG          | UPDATE    | IMS.plxname.PUR.FPPROG    |
| PUR          | FPRGN           | UPDATE    | IMS.plxname.PUR.FPRGN     |
| PUR          | LINE            | UPDATE    | IMS.plxname.PUR.LINE      |
| PUR          | LTERM           | UPDATE    | IMS.plxname.PUR.LTERM     |
| PUR          | MSNAME          | UPDATE    | IMS.plxname.PUR.MSNAME    |
| PUR          | TRAN            | UPDATE    | IMS.plxname.PUR.TRAN      |
| QRY          | AREA            | READ      | IMS.plxname.QRY.AREA      |
| QRY          | DB              | READ      | IMS.plxname.QRY.DB        |
| QRY          | DBDESC          | READ      | IMS.plxname.QRY.DBDESC    |
| QRY          | IMS             | READ      | IMS.plxname.QRY.IMS       |
| QRY          | IMSPLEX         | READ      | IMS.plxname.QRY.IMSPLEX   |
| QRY          | LE              | READ      | IMS.plxname.QRY.LE        |
| QRY          | LTERM           | READ      | IMS.plxname.QRY.LTERM     |
| QRY          | MEMBER          | READ      | IMS.plxname.QRY.MEMBER    |
| QRY          | MSLINK          | READ      | IMS.plxname.QRY.MSLINK    |
| QRY          | MSNAME          | READ      | IMS.plxname.QRY.MSNAME    |
| QRY          | MSLINK          | READ      | IMS.plxname.QRY.MSPLINK   |
| QRY          | NODE            | READ      | IMS.plxname.QRY.NODE      |
| QRY          | OLC             | READ      | IMS.plxname.QRY.OLC       |
| QRY          | OLREORG         | READ      | IMS.plxname.QRY.OLREORG   |
| QRY          | PGM             | READ      | IMS.plxname.QRY.PGM       |
| QRY          | PGMDESC         | READ      | IMS.plxname.QRY.PGMDESC   |
| QRY          | RTC             | READ      | IMS.plxname.QRY.RTC       |
| QRY          | RTCDESC         | READ      | IMS.plxname.QRY.RTCDESC   |
| QRY          | STRUCTURE       | READ      | IMS.plxname.QRY.STRUCTURE |
| QRY          | TRAN            | READ      | IMS.plxname.QRY.TRAN      |
| QRY          | TRANDESC        | READ      | IMS.plxname.QRY.TRANDESC  |

Table 10. Resource names and RACF authority for IMS commands (continued)

| Command verb | Command keyword | Authority | Resource name             |
|--------------|-----------------|-----------|---------------------------|
| QRY          | USER            | READ      | IMS.plxname.QRY.USER      |
| QRY          | USEREXIT        | READ      | IMS.plxname.QRY.USEREXIT  |
| QRY          | USERID          | READ      | IMS.plxname.QRY.USERID    |
| QUE          | LTERM           | UPDATE    | IMS.plxname.QUE.LTERM     |
| QUE          | TRAN            | UPDATE    | IMS.plxname.QUE.TRAN      |
| QUI          | NODE            | UPDATE    | IMS.plxname.QUI.NODE      |
| RDI          | MASTER          | READ      | IMS.plxname.RDI.MASTER    |
| REC          | ADD             | UPDATE    | IMS.plxname.REC.ADD       |
| REC          | REMOVE          | UPDATE    | IMS.plxname.REC.REMOVE    |
| REC          | START           | UPDATE    | IMS.plxname.REC.START     |
| REC          | STOP            | UPDATE    | IMS.plxname.REC.STOP      |
| REC          | TERMINATE       | UPDATE    | IMS.plxname.REC.TERMINATE |
| RMC          |                 | UPDATE    | IMS.plxname.RMC           |
| RMD          |                 | UPDATE    | IMS.plxname.RMD           |
| RMG          |                 | UPDATE    | IMS.plxname.RMG           |
| RMI          |                 | UPDATE    | IMS.plxname.RMI           |
| RML          |                 | READ      | IMS.plxname.RML           |
| RMN          |                 | UPDATE    | IMS.plxname.RMN           |
| RST          | LINE            | UPDATE    | IMS.plxname.RST.LINE      |
| RST          | LINK            | UPDATE    | IMS.plxname.RST.LINK      |
| RST          | MSPLINK         | UPDATE    | IMS.plxname.RST.MSPLINK   |
| RST          | NODE            | UPDATE    | IMS.plxname.RST.NODE      |
| RST          | USER            | UPDATE    | IMS.plxname.RST.USER      |
| RTA          | DUMPQ           | UPDATE    | IMS.plxname.RTA.DUMPQ     |
| RTA          | FREEZE          | UPDATE    | IMS.plxname.RTA.FREEZE    |
| RTA          | UNPLAN          | UPDATE    | IMS.plxname.RTA.UNPLAN    |
| SEC          | APPC            | UPDATE    | IMS.plxname.SEC.APPC      |
| SEC          | OTMA            | UPDATE    | IMS.plxname.SEC.OTMA      |
| SMC          | MASTER          | UPDATE    | IMS.plxname.SMC.MASTER    |
| SMC          | TERMINAL        | UPDATE    | IMS.plxname.SMC.TERMINAL  |
| STA          | APPC            | UPDATE    | IMS.plxname.STA.APPC      |
| STA          | AREA            | UPDATE    | IMS.plxname.STA.AREA      |
| STA          | AUTOARCH        | UPDATE    | IMS.plxname.STA.AUTOARCH  |
| STA          | CLASS           | UPDATE    | IMS.plxname.STA.CLASS     |
| STA          | DB              | UPDATE    | IMS.plxname.STA.DB        |
| STA          | DATAGRP         | UPDATE    | IMS.plxname.STA.DATAGRP   |
| STA          | DC              | UPDATE    | IMS.plxname.STA.DC        |
| STA          | ISOLOG          | UPDATE    | IMS.plxname.STA.ISOLOG    |
| STA          | LINE            | UPDATE    | IMS.plxname.STA.LINE      |
| STA          | LTERM           | UPDATE    | IMS.plxname.STA.LTERM     |

Table 10. Resource names and RACF authority for IMS commands (continued)

| Command verb | Command keyword | Authority | Resource name            |
|--------------|-----------------|-----------|--------------------------|
| STA          | LU              | UPDATE    | IMS.plxname.STA.LU       |
| STA          | MADSIOT         | UPDATE    | IMS.plxname.STA.MADSIOT  |
| STA          | MSNAME          | UPDATE    | IMS.plxname.STA.MSNAME   |
| STA          | NODE            | UPDATE    | IMS.plxname.STA.NODE     |
| STA          | OLDS            | UPDATE    | IMS.plxname.STA.OLDS     |
| STA          | OTMA            | UPDATE    | IMS.plxname.STA.OTMA     |
| STA          | PGM             | UPDATE    | IMS.plxname.STA.PGM      |
| STA          | REGION          | UPDATE    | IMS.plxname.STA.REGION   |
| STA          | RTC             | UPDATE    | IMS.plxname.STA.RTC      |
| STA          | SB              | UPDATE    | IMS.plxname.STA.SB       |
| STA          | SERVGRP         | UPDATE    | IMS.plxname.STA.SERVGRP  |
| STA          | SUBSYS          | UPDATE    | IMS.plxname.STA.SUBSYS   |
| STA          | SURV            | UPDATE    | IMS.plxname.STA.SURV     |
| STA          | THREAD          | UPDATE    | IMS.plxname.STA.THREAD   |
| STA          | TMEM            | UPDATE    | IMS.plxname.STA.TMEM     |
| STA          | TRAN            | UPDATE    | IMS.plxname.STA.TRAN     |
| STA          | TRKARCH         | UPDATE    | IMS.plxname.STA.TRKARCH  |
| STA          | USER            | UPDATE    | IMS.plxname.STA.USER     |
| STA          | VGR             | UPDATE    | IMS.plxname.STA.VGR      |
| STA          | WADS            | UPDATE    | IMS.plxname.STA.WADS     |
| STA          | XRCTrack        | UPDATE    | IMS.plxname.STA.XRCTrack |
| STO          | ADS             | UPDATE    | IMS.plxname.STO.ADS      |
| STO          | APPC            | UPDATE    | IMS.plxname.STO.APPC     |
| STO          | AREA            | UPDATE    | IMS.plxname.STO.AREA     |
| STO          | AUTOARCH        | UPDATE    | IMS.plxname.STO.AUTOARCH |
| STO          | BACKUP          | UPDATE    | IMS.plxname.STO.BACKUP   |
| STO          | CLASS           | UPDATE    | IMS.plxname.STO.CLASS    |
| STO          | DB              | UPDATE    | IMS.plxname.STO.DB       |
| STO          | DATAGRP         | UPDATE    | IMS.plxname.STO.DATAGRP  |
| STO          | DC              | UPDATE    | IMS.plxname.STO.DC       |
| STO          | LINE            | UPDATE    | IMS.plxname.STO.LINE     |
| STO          | LTERM           | UPDATE    | IMS.plxname.STO.LTERM    |
| STO          | LU              | UPDATE    | IMS.plxname.STO.LU       |
| STO          | MADSIOT         | UPDATE    | IMS.plxname.STO.MADSIOT  |
| STO          | MSNAME          | UPDATE    | IMS.plxname.STO.MSNAME   |
| STO          | NODE            | UPDATE    | IMS.plxname.STO.NODE     |
| STO          | OLDS            | UPDATE    | IMS.plxname.STO.OLDS     |
| STO          | OTMA            | UPDATE    | IMS.plxname.STO.OTMA     |
| STO          | PGM             | UPDATE    | IMS.plxname.STO.PGM      |
| STO          | REGION          | UPDATE    | IMS.plxname.STO.REGION   |

Table 10. Resource names and RACF authority for IMS commands (continued)

| Command verb | Command keyword | Authority | Resource name            |
|--------------|-----------------|-----------|--------------------------|
| STO          | RTC             | UPDATE    | IMS.plxname.STO.RTC      |
| STO          | SB              | UPDATE    | IMS.plxname.STO.SB       |
| STO          | SERVGRP         | UPDATE    | IMS.plxname.STO.SERVGRP  |
| STO          | SUBSYS          | UPDATE    | IMS.plxname.STO.SUBSYS   |
| STO          | SURV            | UPDATE    | IMS.plxname.STO.SURV     |
| STO          | THREAD          | UPDATE    | IMS.plxname.STO.THREAD   |
| STO          | TMEM            | UPDATE    | IMS.plxname.STO.TMEM     |
| STO          | TRAN            | UPDATE    | IMS.plxname.STO.TRAN     |
| STO          | USER            | UPDATE    | IMS.plxname.STO.USER     |
| STO          | VGR             | UPDATE    | IMS.plxname.STO.VGR      |
| STO          | WADS            | UPDATE    | IMS.plxname.STO.WADS     |
| STO          | XRCTrack        | UPDATE    | IMS.plxname.STO.XRCTrack |
| SWI          | OLDS            | UPDATE    | IMS.plxname.SWI.OLDS     |
| SWI          | SYSTEM          | UPDATE    | IMS.plxname.SWI.SYSTEM   |
| SWI          | WADS            | UPDATE    | IMS.plxname.SWI.WADS     |
| TERM         | OLC             | UPDATE    | IMS.plxname.TERM.OLC     |
| TERM         | OLREORG         | UPDATE    | IMS.plxname.TERM.OLREORG |
| TES          | MFS             | UPDATE    | IMS.plxname.TES.MFS      |
| TRA          | SET             | UPDATE    | IMS.plxname.TRA.SET      |
| UNL          | DB              | UPDATE    | IMS.plxname.UNL.DB       |
| UNL          | PGM             | UPDATE    | IMS.plxname.UNL.PGM      |
| UNL          | SYSTEM          | UPDATE    | IMS.plxname.UNL.SYSTEM   |
| UNL          | TRAN            | UPDATE    | IMS.plxname.UNL.TRAN     |
| UPD          | AREA            | UPDATE    | IMS.plxname.UPD.AREA     |
| UPD          | DATAGRP         | UPDATE    | IMS.plxname.UPD.DATAGRP  |
| UPD          | DB              | UPDATE    | IMS.plxname.UPD.DB       |
| UPD          | DBDESC          | UPDATE    | IMS.plxname.UPD.DBDESC   |
| UPD          | IMS             | UPDATE    | IMS.plxname.UPD.IMS      |
| UPD          | LE              | UPDATE    | IMS.plxname.UPD.LE       |
| UPD          | MSLINK          | UPDATE    | IMS.plxname.UPD.MSLINK   |
| UPD          | MSNAME          | UPDATE    | IMS.plxname.UPD.MSNAME   |
| UPD          | MSPLINK         | UPDATE    | IMS.plxname.UPD.MSPLINK  |
| UPD          | OLREORG         | UPDATE    | IMS.plxname.UPD.OLREORG  |
| UPD          | PGM             | UPDATE    | IMS.plxname.UPD.PGM      |
| UPD          | PGMDESC         | UPDATE    | IMS.plxname.UPD.PGMDESC  |
| UPD          | RTC             | UPDATE    | IMS.plxname.UPD.RTC      |
| UPD          | RTCDESC         | UPDATE    | IMS.plxname.UPD.RTCDESC  |
| UPD          | TRAN            | UPDATE    | IMS.plxname.UPD.TRAN     |
| UPD          | TRANDESC        | UPDATE    | IMS.plxname.UPD.TRANDESC |
| VUN          | AREA            | UPDATE    | IMS.plxname.VUN.AREA     |

**Related concepts:**

 [Commands for IMS operations tasks \(Operations and Automation\)](#)

**Related reference:**

 [Equivalent IMS type-1 and type-2 commands \(Commands\)](#)

## CSL OM Security user exit routine

OM invokes the optional OM Security user-defined exit routine during command processing. The exit routine performs user-defined security checking when you specify with A or E for the CMDSEC parameter in the OM procedure. The OM Security user-defined exit routine is given control after the OM Input exit routine.

The CSL OM Security user exit routine is defined as TYPE=SECURITY in the EXITDEF statement in the BPE exit list PROCLIB member data set. You can specify one or more user-defined exit routines of this type. When this exit is invoked, all user-defined exit routines of this type are driven in the order specified by the EXITS=*keyword*.

**Related concepts:**

 [Defining the Common Service Layer using IMS PROCLIB data set members \(System Definition\)](#)

**Related reference:**

 [CSL OM Security user exit \(Exit Routines\)](#)

---

## Chapter 10. CSL RM administration

The tasks associated with administering the Resource Manager (RM) are typically performed by the system administrator or system operator.

---

### Information managed by the CSL RM

You can use the Resource Manager (RM) to create, update, query, or delete resources in the resource structure.

Examples of resources whose information is stored on a resource structure include:

- Global resource status information for databases, DEDB areas, and transactions
- Message destination resource information; for example, APPC descriptors, VTAM LTERMS, MSNAMEs, VTAM terminal nodes, users, and user IDs
- Global information for the OLCSTAT data set; for example, current online change libraries
- IMSplex-wide parameters
- Scheduled serial programs

RM maintains the following information about resources on a resource structure:

#### **Resource Name**

A client-defined resource name that is 11 bytes in length.

#### **Name Type**

A resource attribute that ensures that all resources within the same type have unique names. RM enforces this rule. Resources within a name type can have different resource types.

#### **Resource Type**

A client-defined resource attribute that allows CQS to physically group resources on a coupling facility list structure. RM supports up to 255 resource types.

RM uses resource type 253 to store its global information. Resource type 253 hashes to the same physical list headers as resource types 11, 22, 33, 44,...253. A client can define resource types that map to the same physical list headers as RM's global information.

#### **Resource Version**

The number of times that the resource has been updated. RM uses the version to serialize updates on the resource structure.

#### **Resource Owner**

A resource attribute that signifies the owner of the resource. This attribute is optional.

If the owner is identified, the client is responsible for enforcing a single owner of a resource. For example, IMS can set owners for the following resources to enforce a single active instance of the resources: nodes, lterms, users, and user IDs. A user or user ID that is signed on to one IMS in the IMSplex cannot be signed on to another IMS in the IMSplex at the same time. An lterm that is active on one IMS in the IMSplex cannot be active on another IMS in the IMSplex at the same time.



### Resource Data

A resource attribute that contains any additional data about the resource.

To display information about the resource structure managed by RM, use the QUERY STRUCTURE command. More information on this command is in *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

## Maintaining global information for databases, DEDB areas, and transactions

IMS uses the RM resource structure to maintain command status for databases, DEDB areas, and transactions across an IMSplex. This global status enables all IMS systems to view databases, DEDB areas, and transactions in an IMSplex as single databases, areas, and transactions.

Global status for databases, DEDB areas, and transactions is not maintained or obtained on an RSR tracker system. If an IMS system restarts, it uses any global command status that is set in RM and applies that status to its resources, if necessary.

The status of global resources is maintained for most databases whether or not they are registered to DBRC, for DEDB areas, and for transactions. Database status is not maintained for shared secondary index databases and MSDBs.

If an error occurs when RM obtains database or DEDB area status, the global status is not maintained for that resource, and message DFS3308E is issued. DFS3308E lists the RM error return and reason code, followed by message DFS3500I, which indicates that global status for database or DEDB area resources is disabled. All subsequent DB or AREA commands with SCOPE(ALL) or GLOBAL prevent RM status from being updated. After the error is corrected, you can issue UPD IMS SET to enable global status again.

After an online change is committed, global status is obtained for all databases, DEDB areas, or transactions being added. Global command status for CREATE DB and CREATE TRAN commands is also obtained from RM. The databases, DEDB areas, or transactions cannot be used until global status for the resources added is obtained and processed. If an error occurs obtaining or processing the global status for the added resources, the added resources are marked in error.

To maintain global status for databases, DEDB areas, and transactions in RM, specify the appropriate PLEXPARM value during IMS initialization in either the DFSDFxxx or DFSCGxxx IMS PROCLIB member data set.

The status of global resources is set by the first IMS system that starts in an IMSplex, and the values are obtained from the PLEXPARM parameters on that system. By maintaining this information globally using RM, you can, for example, stop a database globally, and any IMS system that joins the IMSplex recognizes that the database is stopped. Similarly, an IMS system joining the IMSplex can be prevented from accessing or updating a database that is in use by an offline process.

Only the last significant global status is maintained in the RM structure for the databases, DEDB areas and transaction resources. If multiple update commands are issued for a resource, only the most recent command status is maintained and available to be queried as the global status. For example, if an UPD TRAN NAME(TRANA) STOP(SCHD) command is issued, the global status of STOSCHD

is maintained for transaction TRANA. If, an UPDATE TRAN NAME(TRANA) STOP(Q) command is then issued, the global status of STOQ is maintained for transaction TRANA. To maintain both the STOSCHD and STOQ status, an UPDATE TRAN NAME(TRANA) STOP(Q,SHCD) command must be issued.

You can also change the PLEXPARMs dynamically using the UPD IMS SET (PLEXPARM()) command. See *IMS Version 11 Commands, Volume 2: IMS Commands N-V* for information about the UPD IMS command.

RM works with CQS to access the resource structure for the client. RM issues the requests to query and maintain the resources, and it notifies the client if there is a resource structure change that affects the client.

The IMS systems in the IMSplex still contain the resource definitions. RM does not ensure resource definition consistency across the IMSplex. You can use global online change to update the resource definitions in the IMS systems.

The following table shows how IMS obtains and updates global status for resources, based on how IMS is started:

*Table 11. How IMS obtains and updates global status for resources*

| Type of start                                       | How status is obtained and updated                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IMS initialization                                  | IMS reads the PLEXPARM values from RM to determine whether the RM value takes precedence, or the IMS definition value should be used. IMS updates the PLEXPARM values in RM for other IMS systems that subsequently come up. If RM global status is maintained, IMS obtains global status of databases, DEDB areas, and transactions from RM's resource structure. |
| Cold start or /ERE COLDBASE or /ERE COLDSYS command | Global status of databases, DEDB areas, and transactions is obtained from RM's resource structure.                                                                                                                                                                                                                                                                 |
| Warm start and emergency restart                    | Local status from log records is applied first. After log record processing completes, if IMS determines that the global status that it read from RM was updated while the IMS system was down, the RM global status that was read from RM is applied to the resource.                                                                                             |

## Information stored for databases resources

The information that RM maintains for databases includes:

- Database names
- Database types
- Database access type
- HALDB master name if the database is a HALDB partition
- Database status

The commands and system actions that cause global status to be updated for databases include:

- /DBDUMP DB with the GLOBAL keyword
- /DBRECOVERY DB with the GLOBAL keyword
- /START DB with the GLOBAL keyword
- /STOP DB with the GLOBAL keyword

- UPDATE DB with any of the following keywords specified: START(ACCESS) SET(ACCTYP) STOP(ACCESS | SCHD | UPDATES) SET(LOCK(ON|OFF)) and SCOPE=ALL

Global database information is deleted from the RM structure using the UPD IMS SET(PLEXPARM(GSTSDB(N))) command.

### Information stored for DEDB area resources

The information that RM maintains for DEDB areas includes:

- DEDB names
- DEDB area access
- DEDB area names
- DEDB area status

The commands and system actions that cause global status to be updated for DEDB areas include:

- /DBRECOVERY AREA with the GLOBAL keyword
- /START AREA with the GLOBAL keyword
- /STOP AREA with the GLOBAL keyword
- UPDATE AREA with any of the following keywords specified START(ACCESS) SET(ACCTYPE) STOP(ACCESS | SCHD) and SCOPE=ALL

Global DEDB area information is deleted from the RM structure using the UPD IMS SET(PLEXPARM(GSTSAREA(N))) command.

### Information stored for transaction resources

The information that RM maintains for transactions includes the transaction name and the transaction status.

The commands and system actions that cause global status to be updated for transactions include:

- UPD TRAN STOP(Q) SCOPE=ALL
- UPD TRAN STOP(SCHD) SCOPE=ALL
- UPD TRAN START(Q) SCOPE=ALL
- UPD TRAN START(SCHD) SCOPE=ALL

To reset global status to NULL, issue the UPD IMS SET(PLEXPARM(GSTSTRAN=N)) command.

## Maintaining message destination resource information

When the RM uses a resource structure in an IMSplex, you can indicate that IMS shares IMS Transaction Manager resources. This message destination resource information includes terminals, users, user IDs, LTERMs, MSNAMEs, and APPC descriptors.

To maintain message destination resource information, specify STM=YES on the DFSDCxxx member of the IMS PROCLIB data set. See *IMS Version 11 System Definition* for more information about DFSDCxxx.

See Chapter 17, "Planning for Transaction Manager resources in an IMSplex," on page 299 for more information on TM Resources.

## Maintaining global information for the OLCSTAT data set

The OLCSTAT data set name is maintained in the RM resource structure. This ensures that all IMS systems in the IMSplex that can perform global online change use the same OLCSTAT data set name.

The first IMS system that starts in an IMSplex writes its OLCSTAT name to the OLCSTAT data set. IMS systems that subsequently start compare their OLCSTAT name to the OLCSTAT name in RM. If the OLCSTAT names do not match, IMS abends with 2800, code 09.

## Maintaining IMSplex-wide parameters

The PLEXPARM parameters, GSTSAREA=, GSTSDB=, and GSTSTRAN=, define whether or not global status is maintained. Values for these parameters are maintained in the RM resource structure.

This ensures that all IMS systems in the IMSplex have the same values for the global PLEXPARM values. The first IMS system that starts in an IMSplex writes its global PLEXPARM values from its DFSCGxxx or DFSDFxxx PROCLIB member data set to the RM resource structure. IMS systems that subsequently start use the values defined in RM. If the PLEXPARM values in the IMS system starting up do not match the PLEXPARM values in RM, message DFS3425I is issued, indicating that the IMS values are overridden by the RM values.

## How IMS maintains global information for sysplex serialized program management

Sysplex serialized program management allows users in a shared queues environment to prevent application programs that are defined as serial from being scheduled in parallel on another IMS system in an IMSplex.

Information about scheduled serial PSBs is maintained in the Resource Manager (RM) resource structure to ensure that the serial PSB is scheduled in only one IMS across an IMSplex at any point in time.

For an IMS in an IMSplex to schedule a serial PSB, that IMS must successfully create a unique instance for the PSB on the RM resource structure before it completes scheduling. If another IMS finds that the PSB instance is not unique, it discontinues the scheduling process for the PSB.

---

## Resource structure duplexing requirements for CSL RM

CQS does not log resource updates or support a checkpoint of the resource structure. However, CQS supports automatic duplexing of the resource structure for backup in case of structure failure.

You must define the resource structure in the coupling facility resource management (CFRM) policy to automatically be duplexed if you want the resource structure to be recoverable.

---

## How the CSL RM repopulates a resource structure

If the resource structure and its duplex (if applicable) fail, and CQS can allocate a new structure, CQS notifies RM to repopulate the structure. RM repopulates the structure from information in its local control blocks.

RM then issues a directive to its clients to populate the structure. Any RM or IMS resource that existed only on the resource structure is lost. When a new resource structure is allocated, the clients can choose to coordinate the repopulation of the resource structure.

Every active IMS system that receives the RM repopulate directive updates RM with the following information:

- Its own PLEXPARM values.
- Its transactions.
- Its local list of database resources, DEDB area resources, and transaction resources that have global status (if global status is maintained).
- The OLCSTAT name
- Terminals, users (or ISC subpools), LTERMs, user IDs, and APPC descriptor names (if STM=YES is specified in DFSDCxxx during IMS initialization) for VTAM resources only

If all IMS systems are down when RM requests repopulation, no IMS information like the global status for databases, DEDB areas, or transactions is available in the new RM structure.

---

## How z/OS rebuilds a resource structure

Resource structures are defined with system-managed rebuild, so z/OS automatically rebuilds the structure when no CQS is up to build the structure. z/OS cannot rebuild the structure if the structure fails or if z/OS loses connectivity to the structure.

If any CQS is up and rebuild is initiated with the SETXCF START,REBUILD command, then CQS copies the structure. If the structure fails, no structure recovery is initiated because resource structures do not support structure checkpoint.

---

## Chapter 11. CSL SCI administration

This topic describes the administrative tasks associated with SCI.

---

### CSL SCI security

To access the IMSplex, clients must first register with SCI. A client can register with SCI only if the user ID of the address space in which the client is running has the authority to do so. Define the authority for a client address space in the RACF FACILITY class profile for the IMSplex.

The profile names in the RACF FACILITY class must be in the form `CSL.imsplex_name`, where *imsplex\_name* is the name of the IMSplex that is being protected. The *imsplex\_name* is the characters “CSL” followed by the IMSplex name as defined on the IMSPLEX parameter in the CSLSIxxx PROCLIB member data set.

SCI checks security when a client issues a CSLSCREG request to register with SCI. In response to the CSLSCREG request, SCI issues a RACROUTE REQUEST=AUTH call to RACF to see if the client has the authority to register with SCI. RACF checks the user ID of the address space that issued the CSLSCREG request. This user ID must have at least UPDATE authority to register with SCI.

**Important:** The SCI address space registers with itself but does not need security authorization.

**Example:** An IMSplex, PLEX1, has the IMSplex ID of CSLPLEX1. To define a profile that allows only users IMSUSER1 and IMSUSER2 to register with SCI in PLEX1, issue the following RACF commands:

```
RDEFINE FACILITY CSL.CSLPLEX1 UACC(NONE)
PERMIT CSL.CSLPLEX1 CLASS(FACILITY) ID(IMSUSER1) ACCESS(UPDATE)
PERMIT CSL.CSLPLEX1 CLASS(FACILITY) ID(IMSUSER2) ACCESS(UPDATE)
SETROPTS CLASSACT(FACILITY)
```



---

## Chapter 12. CQS administration

This topic describes the system administration tasks associated with using the Common Queue Server.

This topic contains Product-sensitive Programming Interface information.

---

### Recording information necessary for starting CQS

CQS uses the z/OS system logger to record all information necessary for CQS to recover queue structures and restart. CQS writes log records for each coupling facility list structure pair that it uses to a separate log stream. The log stream is shared among all CQS address spaces that share the queue structure.

The z/OS system logger provides a merged log for all CQS address spaces that are sharing queues on a coupling facility list structure.

**Important:** Changes to resource structures are not logged.

For CQS to use a z/OS system log, you must first define the log stream and associated resources to z/OS.

CQS also provides a File Select and Formatting Print utility to print the log records.

For more information on defining the log stream, see *z/OS MVS Setting Up a Sysplex*.

For more information on the File Select and Formatting Print utility, see *IMS Version 11 System Utilities*.

**Related concepts:**

“CQS recovery functions” on page 67

---

### Establishing client connection to CQS during failed client takeover

When the client (such as IMS) supports XRF takeover capability, one client must take over the work for a failed client. The CQS connected to this failed client might not be active. Therefore, during the takeover process, the client taking over work from a failed client must connect to a different CQS and indicate that it is taking over work from another client.

At this point, CQS must perform a process similar to CQS restart, using the log records from the CQS connected to the failed client. Normally, CQS failed client connection restart (warm start) is automatic and you do not need to take any action.

During a failed client connection, CQS reads the log token from the structure for the CQS connected to the failed client.

- If CQS finds the log token, CQS performs warm start processing for the failed client.
- If CQS does not find the log token, CQS issues WTOR CQS0033A. At this point, you can do one of the following:



- Cold start the client connection.
- Reject the client connection request.
- Specify a new log token.

**Recommendation:** If a CQS does not accept a log token during failed client connection restart, cold start the connection. If multiple CQSs are running, one CQS structure checkpoint might purge the log records for another CQS that previously failed and was not restarted.

---

## Authorizing access to CQS

If RACF or another security product is installed at your installation, the security administrator can define profiles that control the ability of clients to connect to and access CQS structures.

### Authorizing CQS registration

When a client issues the CQSREG request to register with CQS, CQS issues a RACROUTE REQUEST=AUTH call to determine whether the client is authorized to register with CQS. RACF checks the user ID of the client that issued the CQSREG request. This user ID must have at least UPDATE authority to register with CQS.

The RACF security administrator can define profiles in the FACILITY class to control registration with CQS. The profile names must be of the form `CQS.cqs_id`, where `cqs_id` is the ID of the CQS that is to be protected. The `cqs_id` value is the subsystem name (SSN) as defined in the CQSIPxxx PROCLIB member data set, followed by the characters CQS. For example, if the SSN is ABC, the `cqs_id` value is ABCCQS.

### RACF commands for authorizing CQS registration

To define a profile for CQS to prevent users other than CQSUSER1 and CQSUSER2 from registering, issue the RACF commands shown in the following example.

```
RDEFINE FACILITY CQS.ABCCQS UACC(NONE)
PERMIT CQS.ABCCQS CLASS(FACILITY) ID(CQSUSER1) ACCESS(UPDATE)
PERMIT CQS.ABCCQS CLASS(FACILITY) ID(CQSUSER2) ACCESS(UPDATE)
SETROPTS CLASSACT(FACILITY)
```

### Authorizing connections to CQS structures

When a client issues the CQSCONN request to connect to a CQS structure, CQS issues a RACROUTE REQUEST=AUTH call to determine whether the client is authorized to access the structure. RACF checks the user ID of the client that issued the CQSCONN request. This user ID must have at least UPDATE authority to connect to the structure through CQS.

The RACF security administrator should define profiles in the FACILITY class to control the connection to CQS structures. The profile names must be of the form `CQSSTR.structure_name`, where `structure_name` is the name of the primary CQS structure that is to be protected. Use the same structure name that you define in the CQSSGxxx and CQSSLxxx PROCLIB member data sets.

The `CQSSTR.structure_name` profiles only control access to the specified structures through CQS; they do not control direct access to the structures using IXL macros. You can provide control over direct structure access by defining RACF profiles of

the form `IXLSTR.structure_name`. If you create such profiles, you must give the user IDs under which you run CQS access to the structures.

**Related reading:** For information on protecting direct access to coupling facility structures, see “Authorizing Coupling Facility Requests” in *z/OS MVS Programming: Sysplex Services Guide*.

For more information on defining structure names, see *IMS Version 11 System Definition*. CQS does not perform a separate check on the overflow structure name, because the primary and overflow structures are considered one unit.

**Example:** To define a profile for a CQS primary structure named `IMSMSGQ01`, and to allow only user `CQSUSER` to connect to it, issue the RACF commands shown in the following example.

```
RDEFINE FACILITY CQSSTR.IMSMSGQ01 UACC(NONE)
PERMIT CQSSTR.IMSMSGQ01 CLASS(FACILITY) ID(CQSUSER) ACCESS(UPDATE)
SETROPTS CLASSACT(FACILITY)
```

If you do not define a profile for a particular CQS structure, the structure is not protected, and any user ID can issue a `CQSCONN` request to access the structure.

---

## Using structure alter for CQS

Structure alter is a z/OS process supported by CQS that can be used to alter the structure size or to redistribute the objects within the structure. CQS supports structure alter for primary queue structures, overflow queue structures, and resource structures. CQS allows you to dynamically change the size of a primary or overflow structure.

To enable structure alter, activate a CFRM policy and define the `INITSIZE` and `SIZE` parameters in this policy. For information on structure size, see *IMS Version 11 System Definition*.

To change the structure size, enter the following XES command:

```
SETXCF START,ALTER,STRNAME=structure,SIZE=size
```

The value of *size* must be within the range of values between `INITSIZE` and `SIZE` in the CFRM policy.

Automatic structure alter is a z/OS function that can automatically alter the structure size or the element to entry ratio when the structure full threshold is reached. CQS supports automatic structure alter for queue structures and resource structures. To enable automatic structure alter, activate a CFRM policy defined with `INITSIZE`, `SIZE`, `ALLOWAUTOALT(YES)`.

**Important:** A structure enabled with automatic structure alter is a candidate to be contracted in size by z/OS, if the coupling facility storage becomes constrained. Be careful when enabling automatic structure alter for queue structures. If z/OS contracts the queue structure size, it might cause the queue structure to go into overflow unnecessarily. To prevent this happening, define the CFRM policy with a `MINSIZE` (minimum size), below which z/OS will not contract the structure.

The `FULLTHRESHOLD` parameter in the CFRM policy structure definition specifies the percent full threshold for the structure. This threshold is the percent full that coupling facility resources such as entries or elements must reach before auto alter processing is triggered. The default value for the `FULLTHRESHOLD`

parameter is 80 percent. The CQSSGxxx PROCLIB member parameter OVFLWMAX specifies the percentage of elements in the structure that must be in use to trigger CQS overflow processing. The default value for OVFLWMAX is 70 percent.

**Recommendation:** If you are using ALLOWAUTOALT(YES), change one or both of these default values so that the OVFLWMAX value is larger than the FULLTHRESHOLD value by at least five percent. This is to give z/OS a chance to alter the entry-to-element ratio and structure size, and potentially avoid CQS overflow processing. Also, setting these parameters to different values reduces the chance of both auto alter and CQS overflow processing occurring at the same time and causing extra overhead.

For more information on automatic Structure Alter, see *z/OS MVS Setting Up a Sysplex*.

---

## Using CQS system checkpoint

CQS system checkpoint, the checkpoint data sets that are used for recovery, applies to a CQS if it manages at least one queue structure. If a CQS manages only a resource structure, system checkpoint does not apply.

At a system checkpoint for recovering CQS information, CQS writes log records that contain restart and recovery information to the CQS log. CQS does not stop activity while the checkpoint is in progress.

CQS performs a system checkpoint in each of the following situations:

- When a client issues a CQSCHKPT FUNC=CHKPTSYS request
- When the number of log records that CQS writes reaches the value specified on the SYSCHKPT= parameter in the CQSSLxxx PROCLIB member data set
- When the client is IMS and you enter the /CQCHKPT SYSTEM command
- When a client RESYNC ends
- When structure checkpoint ends successfully
- At the end of a restart

In addition, CQS takes system checkpoints during significant events, such as a shutdown.

### CQS checkpoint data set

For each structure pair, CQS maintains a checkpoint data set. CQS writes to its checkpoint data set and uses it during restart.

The checkpoint data set is dynamically allocated during CQS initialization. You define the checkpoint data set DSNAME for a structure using the CHKPTDSN= parameter in PROCLIB member data set CQSSLxxx.

### How CQS restarts after system checkpoint

During CQS restart, CQS reads the log records from the last system checkpoint and restores the environment for committed data objects and backs out uncommitted data objects on queue structures. The frequency of system checkpoint affects this restart. CQS must read more log records when checkpoints are infrequent than when the checkpoints occur more often. Because the CQS log is shared by multiple

CQSs, CQS restart time is affected by the number of log records written by the multiple CQSs, not just the CQS that is being restarted.

CQS takes an initial system checkpoint at the end of a restart.

**Related concepts:**

“CQS system checkpoint data set” on page 98

---

## Using CQS structure checkpoint

Structure checkpoint takes a snapshot of the shared queues on a queue structure and writes the data to the structure recovery data set (SRDS) so that CQS can recover the queues after a structure failure. Structure checkpoint processing copies all recoverable data objects from a structure pair to a SRDS.

For nonrecoverable data objects, the queue name, and UOW are copied, but not the actual data object. The client specifies whether or not a data object is recoverable when the CQSPUT FUNC=PUT request is issued to insert the data object onto the shared queues. For example, when IMS is the client, all data objects are marked as recoverable, except for Fast Path input messages.

**Important:** Structure checkpoint is not supported for resource structures. It supports queue structures only.

When it performs the copy operation, CQS stops all activity against the structure to ensure that the structure does not change while the checkpoint is being taken. If CQS receives a request to process work when a structure checkpoint is in progress, the request is held until after the structure checkpoint is complete.

**Recommendation:** Because no other work for a structure can be processed while CQS is taking a checkpoint, consider processing structure checkpoints during non-peak hours.

After all shared queues are copied to the SRDS, each CQS performs a system checkpoint to ensure its restart checkpoint has a time stamp that is more recent than the current structure checkpoint. The structure checkpoint process then deletes all log records that are not needed for structure recovery, allowing the logger to reclaim space in the CQS log and preventing the log from becoming full. After log records are deleted, CQS cannot access these log records and, therefore, cannot use these records for structure recovery or CQS restart. If only one SRDS contains valid structure checkpoint data, all log records that were written prior to that structure checkpoint are deleted. If both SRDSs contain valid structure checkpoint data, all log records that were written prior to the oldest structure checkpoint are deleted.

If a CQS was not active at the time of a structure checkpoint, it cannot initiate a system checkpoint, meaning that its restart checkpoint is older than at least one structure checkpoint. If both SRDSs contain valid structure checkpoint data, no problem exists (because the CQS restart checkpoint is still more recent than the oldest structure checkpoint, so its restart log records are not deleted). However, if this is the first or only valid structure checkpoint, or a CQS was down across two structure checkpoints, the log records needed for that CQS to restart are deleted. In this case, that CQS might need a cold start to restart.

**Recommendation:** Initiate a structure checkpoint after a structure cold start, or anytime the SRDSs are deleted and redefined. The structure checkpoint should

start after all CQSs that share the structure are started. This provides the SRDS with an initial structure checkpoint. Also, to update the snapshot of the shared queues and to periodically delete log records, initiate a structure checkpoint at regular intervals.

When a structure recovery is required, the SRDS and the CQS log are used to recover the shared queues. CQS first repopulates the new structure from the SRDS. CQS then reads all log records from the time the structure checkpoint completed. The length of time to read the log records is dependent on how many log records are in the log. More frequent structure checkpoints reduce the number of log records that must be read during a structure recovery. Deleting the log records also helps prevent the log from becoming full. When a log stream becomes full, CQS deletes all log records older than the oldest structure checkpoint or CQS system checkpoint. CQS then takes a structure checkpoint.

CQS performs structure checkpoints in each of the following situations:

- When the z/OS log becomes full or approaches full.
- After a successful structure recovery.
- After a successful overflow threshold process.
- When a client issues the CQSCHKPT FUNC=CHKPTSTR request.
- When the client is IMS, and you enter the /CQCHKPT SHAREDQ command.
- During CQS normal termination when the client requests it on the CQSDISC request. When the client is IMS, you can request a structure checkpoint at CQS termination by entering the /CQSET SHUTDOWN SHAREDQ command. IMS then passes this request to CQS when IMS terminates normally with a /CHECKPOINT FREEZE|DUMPQ|PURGE command.

**Related concepts:**

“CQS recovery functions” on page 67

---

## Preventing CQS structure full

You should manage structure usage to avoid a structure full condition. If a resource structure or queue structure becomes full, CQS issues message CQS0205E.

There are various ways to prevent a structure full condition:

- CQS structure overflow function for queue structures
- z/OS structure full monitoring capability, used with CQS, for queue structures and resource structures

Use these mechanisms to warn when a structure full condition is approaching and to take action to prevent a full structure.

**Related concepts:**

“CQS structure functions” on page 66

## CQS structure overflow function

CQS provides a structure overflow function that automatically warns you when a queue structure is approaching full and takes action to prevent a full structure. When the usage of a structure reaches the *overflow threshold*, CQS attempts to make the structure larger by initiating a structure alter. If the alter fails, CQS either allocates an overflow structure and moves data objects associated with selected queues to the overflow structure (if you define an overflow structure) or rejects

data objects from being put on the selected queues. CQS stops all activity against the structure while the queues are being selected.

**Definition:** The *overflow threshold* is the percentage of the primary structure that must be in use before CQS goes into overflow mode. The default overflow threshold is 70%, but you can change the default by defining the OVFLWMAX parameter in the CQSSGxxx member of the PROCLIB data set.

**Important:** Structure overflow is not supported for resource structures.

If CQS does not succeed in altering a structure's size, the structure goes into *overflow mode*. In overflow mode, CQS selects queues using the most space on the structure as candidates for *overflow processing*. CQS stops selecting queues when enough queues have been selected to cause the primary structure usage to fall 20% below the overflow threshold. Activity against the structure is temporarily stopped while queues are being selected for overflow. CQS drives the Queue Overflow User-Supplied exit routine with the candidate queue names, which the exit then approves or rejects for overflow processing. Queues that get approved are placed into overflow mode. If an overflow structure is defined, CQS allocates the overflow structure and moves the data objects associated with the approved queues to the overflow structure. If an overflow structure is not defined, CQS rejects CQSPUT requests for the approved queues. Overflow structures can be defined in the CQSSGxxx member of the PROCLIB data set, using the OVFLWSTR parameter.

If an overflow structure is allocated, a queue remains in overflow mode until CQS detects that there are no more entries in the queue. (The queues are scanned every 15 minutes.) After the empty queue is detected, it is removed from overflow mode and returned to normal processing. Because the scan is periodic, empty queues might remain in overflow mode temporarily. When all queues have been removed, CQS exits overflow mode. If no overflow structure is allocated, CQS exits overflow mode when the primary structure usage drops 20% below the overflow threshold.

**Remember:** Large messages (messages that exceed the size of a queue manager long message record) are put on the shared queue structure in multiple parts. The first part of a message is placed on the transaction or LTERM ready queue. The second and subsequent parts are placed on a secondary queue called the *staging queue*. When a transaction or LTERM queue is selected for overflow processing, only the first part of any large messages on that queue are moved to the overflow structure. Any parts of a message that are on the staging queue remain on the primary structure. The staging queue is not eligible for overflow processing. You can use the LGMSGSZ parameter to change the size of the queue manager long message record.

## CQS structure full monitoring

The z/OS structure full monitoring capability can be used for queue structures and resource structures to warn you when a structure is approaching full and to prevent a full structure.

If structure full monitoring is enabled, z/OS monitors structure usage. When the number of entries or elements in use reaches the structure full threshold, z/OS issues a highlighted IXC585E message to warn the system programmer that a structure full condition is imminent. If automatic altering is enabled, z/OS automatically initiates a structure alter to increase the structure size or change the element-to-entry ratio.



Structure full monitoring is automatically enabled with a default threshold of 80%. Define a different threshold with the CFRM policy FULLTHRESHOLD parameter. Define the CFRM policy with FULLTHRESHOLD(0) to disable structure full monitoring. When the structure usage goes below the threshold, z/OS issues an IXC586I message.

The following command displays the structure full threshold that is in effect:

```
D XCF,STRUCTURE,STRNAME=strname
```

In the following example, the command

D XCF,STRUCTURE,STRNAME=IMSRSRC01 is issued and the structure full threshold is 80%.

```
STRNAME: IMSR SRC01
STATUS: NOT ALLOCATED
POLICY SIZE      : 4096 K
POLICY INITSIZE  : N/A
FULLTHRESHOLD    : 80
REBUILD PERCENT  : N/A
DUPLEX           : DISABLED
PREFERENCE LIST  : LF03
ENFORCEORDER     : NO
EXCLUSION LIST   IS EMPTY
```

The following example shows the IXC585E message, indicating the structure is full because all of the entries are in use:

```
*IXC585E STRUCTURE IMSR SRC01 IN COUPLING FACILITY LF03, 725
PHYSICAL STRUCTURE VERSION B4704775 92D95302,
IS AT OR ABOVE STRUCTURE FULL MONITORING THRESHOLD OF 80%.
ENTRIES: IN USE:      4874 TOTAL:      4874, 100% FULL
ELEMENTS: IN USE:      19 TOTAL:      4872,   0% FULL
```

The following example shows the IXC586I message:

```
IXC586I STRUCTURE IMSR SRC01 IN COUPLING FACILITY LF03, 772
PHYSICAL STRUCTURE VERSION B4704775 92D95302,
IS NOW BELOW STRUCTURE FULL MONITORING THRESHOLD.
```

**Related reading:** For more details on structure full monitoring and the FULLTHRESHOLD and ALLOWAUTOALT keywords in the CFRM policy, see *z/OS MVS Setting Up a Sysplex*.

## Using structure full monitoring with CQS structure overflow

You can use the structure full monitoring function (a z/OS function) with the structure overflow function (a CQS function) for queue structures.

The overflow threshold is a value defined to CQS. The structure full threshold is a value defined to z/OS. If the overflow threshold is close to the structure full threshold and automatic altering is enabled, CQS and z/OS might both try to initiate a structure alter at the same time to prevent the structure from becoming full.

If a CQS-initiated structure alter is in progress when z/OS detects the structure full threshold has been reached, z/OS stops the CQS-initiated structure alter and initiates its own structure alter. When CQS detects that its structure alter has failed, CQS goes into overflow mode, even if the z/OS-initiated structure alter reduces the structure usage below the overflow threshold.

**Recommendation:** Consider your structure full threshold when deciding what overflow threshold to define, so that you control when a structure goes into overflow mode. If you use structure full threshold, define it to be lower than the overflow threshold to avoid going into overflow mode unnecessarily. If the structure full threshold is lower than the overflow threshold, z/OS can attempt structure full threshold processing before the structure goes into overflow mode.

**Related reading:**

- For detailed information about the CQSSGxxx member of the PROCLIB data set, see *IMS Version 11 System Definition*.
- For detailed information about the Queue Overflow User-Supplied exit routine, see *IMS Version 11 Exit Routines*.
- For detailed information about the CQSPUT request, see *IMS Version 11 System Programming APIs*.

---

## Rebuilding structures in CQS

Structure rebuild is a z/OS process that allows another instance of a structure to be allocated with the same name and data reconstructed from the initial structure instance. z/OS supports system-managed rebuild, CQS-managed rebuild, and structure duplexing.

CQS supports system-managed rebuild and CQS-managed rebuild for queue structures and resource structures. Note that CQS stops all activity against the structure during structure rebuild.

**Related concepts:**

“CQS structure functions” on page 66

### z/OS system-managed rebuild and CQS

System-managed rebuild is a z/OS process by which z/OS rebuilds the structure. z/OS copies the structure contents to a new structure. System-managed rebuild is supported for queue structures and resource structures. System-managed rebuild is only done if no CQS is up. If a CQS is up, the CQS performs a user-managed rebuild and does the structure copy.

Use system-managed rebuild primarily for planned reconfiguration. If the rebuild is initiated with the SETXCF START,REBUILD command and no CQS is available to perform the structure copy, z/OS performs the structure copy.

**Restriction:** System-managed rebuild does not address coupling facility failures, structure failures, or loss of connectivity. CQS-managed rebuild is required to handle such failures.

To enable a structure for system-managed rebuild, add the following parameter to your CFRM couple data set utility job, then run the job control language (JCL) to format the CFRM couple data set with system-managed rebuild capability:

```
ITEM NAME(SMREBLD) NUMBER(1)
```

### CQS-managed rebuild

CQS-managed rebuild is a process by which CQS manages structure rebuild. CQS supports two variations of CQS-managed rebuild: structure copy and structure recovery.



Structure copy copies the contents of the structure to another structure, for a planned reconfiguration or connectivity loss. Structure copy can also be used to activate new CFRM policy attributes. Structure recovery recovers a structure from the SRDS and the z/OS log after a structure failure.

If one CQS loses connectivity to a structure and another CQS still has connectivity to that structure, CQS manages the structure rebuild and performs a structure copy. If all CQSs lose connectivity to a resource structure, structure recovery is attempted, but fails because structure recovery is not supported for resource structures or for nonrecoverable queue structures (RECOVERABLE=NO).

If a coupling facility or queue structure fails, CQS performs a structure recovery.

If a resource structure fails, it is lost and structure rebuild is not performed. CQS is not able to perform structure recovery because resource structures do not support checkpoint and logging. CQS clients can repopulate the failed resource structure. CQS attempts to allocate a new resource structure.

If a new structure is successfully allocated, CQS drives the client structure event exit with the repopulate structure event. The CQS client or clients must then repopulate the structure. If a new structure is not successfully allocated, CQS drives the structure exit event with the structure failed event. The structure is not accessible for repopulation. Correct the environmental problem that caused the structure allocate to fail so that the structure can be allocated and repopulated.

## Initiating structure rebuild with z/OS and CQS

A structure rebuild can be initiated by a z/OS operator, by CQS, or by z/OS.

- A z/OS operator can initiate a structure rebuild to copy or recover queues using the following command:  
`SETXCF START,REBUILD,STRNAME=strname,LOCATION=NORMAL/OTHER`
- CQS initiates a structure rebuild if, during CQS initialization, it detects an empty structure and a valid SRDS (indicating a valid structure checkpoint in the SRDS). If CQS detects an empty structure and a valid SRDS, it also initiates a structure rebuild during event notification facility (ENF) 35 event processing.
- z/OS initiates a structure rebuild if the rebuild threshold for loss of connectivity is reached. The rebuild threshold for loss of connectivity is defined with the CFRM policy REBUILDPERCENT keyword. The REBUILDPERCENT default is 1. If the system programmer does not define REBUILDPERCENT, z/OS initiates a rebuild if any CQS loses connectivity to the structure.
- If structure copy aborts because of a CQS failure and no other CQS can determine if the failed CQS is the master, then the rebuild starts over as a structure recovery.

## CQS structure repopulation

Structure repopulation is a process by which CQS clients repopulate a failed resource structure. CQS does not support structure recovery for resource structures because CQS does not log or checkpoint resource updates.

If a resource structure and its duplex fail, the CQS clients can repopulate the resource structure. CQS attempts to allocate a new structure. If this allocation is successful, CQS notifies its clients to repopulate. The CQS client or clients must then repopulate the structure. Any resources that were kept only on the resource structure are lost.

If CQS fails to allocate a new structure, CQS notifies the client that the structure failed. If the sysplex environment changes later and CQS is eventually able to allocate a new resource structure, CQS notifies the client to repopulate at that time. Alternately, correct the environmental problem that caused the structure allocate to fail so that the structure can be allocated and repopulated.

CQS does not coordinate resource structure repopulation between CQS clients; clients must synchronize resource structure repopulation if desired. Structure repopulation does not guarantee the restoration of all objects; some objects might be lost.

## CQS structure recovery

The structure recovery function recovers the data objects on a structure from the SRDS and the z/OS logs after a structure failure.

**Important:** Structure recovery is not supported for resource structures.

After a structure failure, the structure might need to be recovered if it is empty or contains only CQS control information. During structure recovery, CQS allocates a structure and repopulates it from either the SRDS (containing valid client data from a previous checkpoint) and the CQS log or the CQS log by itself.

When CQS recovers the structure from a structure checkpoint, it repopulates the structure with the data objects from the structure recovery data set. CQS reads the log starting at the time of the structure checkpoint to update the structure with changes that occurred after the structure checkpoint.

If the primary structure is empty and neither SRDS contains valid structure checkpoint data, CQS determines whether it can use just the CQS log for recovery. If the first log record in the log stream is the Beginning of Log log record, the log stream contains all of the log records required for recovery and CQS can use the log record to complete the structure recovery.

If CQS finds that a previous structure rebuild did not complete successfully, it initiates another rebuild.

If the primary structure contains only CQS control information and the CQS that allocated the structure is not able to determine if a rebuild is necessary, CQS initiates a rebuild if either SRDS is valid or all log records are available.

If neither SRDS is valid and the log records are deleted by a previous structure checkpoint, CQS cannot rebuild the structure. In this case, if rebuild is necessary, CQS issues WTOR CQS0034A to ask you what to do. You can cold start the structure or cancel this CQS.

If no CQS has access to the structure when structure rebuild is initiated, the structure is recovered from the SRDS and the CQS log. Nonrecoverable data objects (such as IMS Fast Path input messages) are lost. Data objects are read from the SRDS and copied into a new structure. CQS then reads the log to bring the structure back to the point of currency. The log contains all the records necessary for structure recovery if no structure checkpoint was ever initiated. In this case, the structure is recovered from just the CQS log.

A client can use the CQSCONN request to specify whether work can be performed while a structure is being rebuilt. While structure recovery is in progress, CQS

stops all activity against the structure. This means that CQS requests are held until the structure recovery is complete. You can allow CQS requests to continue during structure rebuild by specifying WAITRBLD=NO when connecting to the structure with the CQSCONN request. In this case, structure recovery stops structure activity for some time, but the structure becomes available much sooner.

**Related concepts:**

“CQS structure recovery” on page 187

## CQS structure copy

The structure copy function copies all of the data objects (both recoverable and nonrecoverable) from the structure to a new structure for a planned reconfiguration or unplanned activity such as loss of connectivity.

Structure copy can be used to change the location of the structure or any other attribute defined in the CFRM policy, such as SIZE, INITSIZE, and PREFLIST. When a structure rebuild is initiated, at least one CQS must have access to the structure for structure copy to be performed.

## z/OS structure duplexing for CQS

Structure duplexing is an optional z/OS-managed process for failure recovery of queue structures and resource structures. In this process, z/OS creates a duplex copy of a structure in advance of a failure, then maintains the structures in a duplexed state during normal operation.

If a queue structure fails and duplexing is enabled, z/OS switches to the unaffected structure instance. If a queue structure fails and duplexing is not enabled, CQS rebuilds the structure based on data from the most recent checkpoint and z/OS log entries. The advantage of duplexing queue structures in the event of a failure is in avoiding the overhead of a CQS-managed structure rebuild.

If duplexing is enabled and a resource structure fails, z/OS switches to the unaffected structure instance. If duplexing is not enabled and a resource structure fails, the data objects are lost because resource structures do not support checkpoint or logging. CQS repopulates the resource structure with control information. CQS notifies its clients to repopulate the structure. It is up to the clients to repopulate the resource structure if necessary.

**Recommendation:** Enable structure duplexing for resource structures.

If both instances of a structure fail at the same time, structure duplexing does not work and all data objects are lost. If the failed structure is a resource structure, the CQS client must repopulate it. If the failed structure is a queue structure, CQS recovers the structure using structure rebuild.

Structure duplexing is optional. To use it, you must enable the z/OS duplexing function. Perform the following steps to enable this function:

1. Ensure that the sysplex is defined as duplexing capable.
2. Add the following parameter to your CFRM couple data set format utility: ITEM NAME(SMDUPLEX) NUMBER(1)
3. Migrate to an environment in which system-managed duplexing is enabled from a CFRM standpoint. A nondisruptive migration of CFRM couple data sets is required. Only z/OS systems at a level that supports system-managed duplexing are capable of using system-managed CFRM couple data sets that are duplexing-capable. Therefore, take the following steps:

- a. Incrementally migrate all systems in the sysplex that are using CFRM to the z/OS level that supports system-managed duplexing.
- b. Format system-managed duplexing-capable CFRM couple data sets and bring them into use as the primary and alternate CFRM couple data sets for the configuration.

**Important:** After you enable z/OS duplexing, you cannot return to downlevel CFRM couple data sets (ones that are not system-managed duplexing-capable) without disruption. Doing so requires a sysplex-wide IPL of all systems using the system-managed duplexing-capable data sets.

After an uplevel CFRM couple data set is in use in the sysplex, system-managed duplexing can be started and stopped in a nondisruptive manner. To turn this function on or off, even while the CFRM couple data set is in use, modify the CFRM policy DUPLEX parameter or use the SETXCF START/STOP,REBUILD,DUPLEX operator command.

To enable system-managed duplexing for a particular structure, the structure must be defined as duplexing-capable. Defining a structure as duplexing capable also defines it as system-managed rebuild-capable. Add the following parameter to your CFRM active policy:

DUPLEX (ENABLED)

or

DUPLEX(ALLOWED)

If DUPLEX(ENABLED) is defined in the CFRM active policy, the system programmer or z/OS internally can initiate the duplexing rebuild. z/OS triggers the start of duplexing rebuild based on a timer or upon detection of certain events (such as connect, disconnect, and policy change). When CQS initializes and connects to a structure defined with DUPLEX(ENABLED), z/OS starts a duplexing rebuild.

If DUPLEX(ALLOWED) is defined in the CFRM active policy, the duplexing rebuild must be initialized by the system programmer using the following command:

SETXCF START,REBUILD,DUPLEX,STRNAME=strname

**Important:** If you define overflow structures with DUPLEX(ENABLED), IMS initialization allocates the overflow structure and duplexing begins. If IMS initialization determines that the overflow structure is not needed, it deletes it and duplexing terminates. If you want to avoid this unnecessary overhead, during CQS initialization define the overflow structure with DUPLEX(ALLOWED) and initiate duplexing with a SETXCF command when the structure goes into overflow mode.

Once duplexing is established, the structure remains in that state indefinitely. Duplexing can be stopped internally by z/OS if an error occurs (such as link failure, structure failure, and CFRM policy change). The system programmer can explicitly stop duplexing using the following command, where you specify KEEP=OLD to keep the old structure and KEEP=NEW to keep the new structure.

SETXCF STOP,REBUILD,DUPLEX,STRNAME=strname,KEEP=OLD/NEW

Planned reconfiguration (such as a CFRM policy change or taking a coupling facility offline for maintenance) is supported. Structure rebuild is not permitted for a structure that has established duplexing, so the duplexing must be stopped first. Perform the following steps:

1. Stop duplexing.

Stop duplexing and switch the structure to simplex mode by issuing the following command:

```
SETXCF STOP,REBUILD,DUPLEX,STRNAME=strname,KEEP=OLD/NEW
```

2. Reconfigure.

Make the change required for planned reconfiguration.

3. Initiate duplexing rebuild.

Initiate a new duplexing rebuild by issuing the following command:

```
SETXCF START,REBUILD,DUPLEX,STRNAME=strname
```

---

## Chapter 13. IMS service considerations

Proper planning for implementing service to IMS is a vital part of keeping IMS available and resilient.

---

### Types of service SMP/E SYSMODs

IMS provides maintenance packaged in SMP/E format.

IMS maintenance is packaged as one of three types of SMP/E SYSMODs:

- Program Temporary Fixes (PTFs)

Program Temporary Fixes (PTFs) are normally considered preventative service. PTFs contain solutions for problems and are distributed in machine-readable format. The PTF is considered the final solution for a problem for the release of IMS for which it is provided.

For modules that supersede a previous level of a module, the source changes are the cumulative delta source changes for the module. If a PTF has a prerequisite, the source changes included in the PTF are not cumulative, but reflect only the code changed for the PTF.

- Authorized Program Analysis Reports (APARs)

Authorized Program Analysis Report (APAR) fixes are considered corrective service. APARs contain solutions for problems and are distributed in machine-readable format. The APAR is considered an interim solution, or temporary solution, for a problem. The final solution is the corresponding PTF or PTFs created at the end of the APAR process. One APAR can become one or more PTFs.

- USERMODs

IMS provides USERMODs in the following situations:

- As an APAR fixtest, to ensure that the problem reported by an APAR is corrected or to provide relief until the APAR or PTF is available
- As a circumvention to a problem, to provide relief until the final APAR or PTF is available
- As a trap (or specialized code) to obtain additional documentation or information (such as a dump) necessary to analyze and understand a problem

**Related reading:** For more detailed SMP/E information, see *SMP/E Reference*.

---

### Service SYSMOD packaging

The IMS service process typically makes APARs available as soon as they are completed, which is normally a few weeks before the corresponding PTF or PTFs are completed. In situations in which a fix is urgently needed after the APAR is completed, but before the PTF is available, using the APAR might be the best short-term solution.

APARs provided by IMS define as prerequisites (PRE, IFREQ, and so on) only those SYSMODs for which the APAR has code dependencies. The APARs list other APARs, not PTFs, as prerequisites.

PTFs contain as requisites (PRE, IFREQ, and so on) all prior PTFs affecting the same elements. Processing a PTF might require the processing of many additional SYSMODS, while processing an APAR might not. In emergency situations where a problem exists and a solution must be implemented quickly, the APAR might be the best short-term solution as it might require the least amount of change. However, you must always use the PTF as the final fix.

When processing APARs, encountering regression messages from SMP/E is normal. These messages must be analyzed to ensure that no regression will occur. If needed, contact IBM Software Support for assistance.

PTFs supersede (SUP) their corresponding APARs. Therefore, removing the APAR prior to processing the PTF is not required.

**Important:** APAR and USERMOD fixes should not be processed using the SMP/E ACCEPT command. The corresponding PTF or PTFs should be processed as the final fix

USERMODs provided by IMS define as prerequisites (PRE, IFREQ, and so on) only those SYSMODs for which the USERMOD has code dependencies. USERMODs list the corresponding APARs, not PTFs, as prerequisites. In this way, USERMODs are like APARs. Whenever IMS USERMODS are processed by SMP/E, regression messages might be encountered. These messages must be analyzed to ensure that no regression is actually taking place. If needed, contact IBM Software Support for assistance.

USERMODs provided by IMS are not superseded (SUP) by a corresponding APAR or PTF. When the final fix is available, you must RESTORE the USERMOD from the system. The ++HOLD information provided with each USERMOD indicates that you must RESTORE the USERMOD and contains instructions on how to do so.

---

## Obtaining IMS service

The recommended way to obtain IMS service is to order a Recommended Service Upgrade (RSU), which includes all current PTFs that have passed the Consolidated Service Test (CST).

CST is an additional service testing environment for the z/OS platform to test all the current PTFs for participating products on the z/OS platform together so that IBM can recommend PTF service for z/OS and key subsystems together in one RSU.

To order your IMS service, go to ShopzSeries at <https://www14.software.ibm.com/webapp/ShopzSeries/ShopzSeries.jsp>.

One of the ShopzSeries options is to create a custom order that is based on the current contents of your SMP/E Consolidated Software Inventory (CSI). If you use this option, ShopzSeries requires information on existing features, FMIDs, and PTFs in your environment. This information can be generated by using the SMP/E GIMXSID service routine. The GIMXSID service examines the specified target zones and global zone within the specified SMP/E CSI to determine which functions and service you already have. The ShopzSeries web site provides detailed instructions, sample JCLs, and video clips that guide you through the service ordering process.



You can also continue to obtain PTFs from the following channels. However, keep in mind that these PTFs may not have gone through the CST process.

#### **IBM Software Support**

You can request specific PTFs that can be downloaded from IBMLINK, a File Transfer Protocol (FTP) site, or mailed on a cartridge.

#### **Extended Support Offering (ESO)**

ESO tapes are available to licensed users upon request.

#### **Custom Built Product Delivery Offering (CBPDO)**

CBPDO service tapes are created upon customer request.

#### **ServerPac**

ServerPac orders are sent upon request. In addition to products, ServerPac orders also contain PTFs that have been incorporated in the products.

#### **ShopzSeries**

Web-based productivity tool that makes it easier for you to order service. Service is sent to you either through the mail or through the Internet.

---

## **Maintenance recommendations**

Maintaining your IMS systems can be a challenge in today's complex environments. The recommendations outlined in the following topics are designed to assist you in developing a maintenance strategy that works in your environment.

Before you apply any maintenance, you need to determine your readiness. This involves a careful risk assessment. This same principle applies to developing a maintenance strategy.

### **Assessing your readiness to apply maintenance**

Numerous factors are involved in assessing your readiness to apply maintenance, including the quality of the local test environment and the business cycle.

- **Quality of the local test environment**

Several factors affect the quality of the test environment. Some questions to consider are these:

- Are the software products such as RACF, DFP, or z/OS DFSMS used in your test environment in the same way they are used in production and do the levels match what is used in production?
- Do you have a tool, such as Teleprocessing Network Simulator (TPNS), that enables you to perform stress tests to simulate peak production activity?
- To what extent does the hardware used in your test environment match the production environment?
- To what extent does the application software used in your test environment match the production environment?
- Are test results closely monitored?

These and other factors need to be considered to evaluate the quality of the test environment.

- **Business cycle**

You must do everything you can during a critical business cycle, to ensure that IMS remains available. Therefore, avoid implementing maintenance to a production system during a critical business cycle.



## General maintenance recommendations for a production system

After starting with a base implementation service level of an IMS production system, install service levels and PTFs, resolve PEs, and conduct a test cycle.

**Attention:** SYSMODs in APPLY-only status could be regressed by an IMS system definition. See “Preventing regression of SYSMODs in APPLY-only status by an IMS system definition” on page 201 for instructions about preventing this.

**Note:** You can use the `/DIAGNOSE SNAP MODULE(modulename)` command to determine your current maintenance level for a specific module.

If you are starting with a base implementation service level of an IMS production system, take the following actions:

1. Install service levels 6 months behind the current ESO or CBPDO level for a planned migration.
2. Install all HIPER PTFs fixes 3 months behind the planned migration date.
3. Resolve PEs.
4. Conduct a 3-month test cycle prior to implementing on your production system.

## General maintenance recommendations for IMS service distributed after testing cycle begins

IMS maintenance continues to be distributed during the 3-month test cycle.

Include the following maintenance items in the initial production base system implementation:

1. Install fully tested fixes for significant software problems encountered.
2. Install fully tested HIPER SYSMODs that are significant and directly apply to your specific environment.
3. Review the IMS PSP bucket UPGRADE for the IMS release level and SUBSET FMIDs; important IMS product information is continually added to these buckets. You might be prompted to take a variety of actions to keep your IMS system in proper working order.

## Sample implementation of maintenance recommendations

Use the following example for a strategy for implementing a maintenance policy.

This maintenance policy is based on the recommendations outlined in “General maintenance recommendations for IMS service distributed after testing cycle begins” and “General maintenance recommendations for a production system.”

1. Obtain current service.

Use your current service delivery method, or use ShopzSeries. ShopzSeries is IBM's web-based productivity tool that simplifies the ordering of zSeries software products, product upgrades, and system maintenance. Go to [www14.software.ibm.com/webapp/ShopzSeries/ShopzSeries.jsp](http://www14.software.ibm.com/webapp/ShopzSeries/ShopzSeries.jsp) for more information about ShopzSeries.

2. Obtain and receive current enhanced HOLDDATA.

Go to [service.software.ibm.com/holddata/390holddata.html](http://service.software.ibm.com/holddata/390holddata.html) for the most current information regarding enhanced HOLDDATA.

3. Use SMP/E to install the service.

Select RSUyymm SOURCEIDs that are created by CST (Consolidated Service Test). Go to <http://www.ibm.com/servers/eserver/zseries/zos/servicetst/> for additional information regarding CST. Then resolve both PEs and System HOLDS.

4. Obtain and receive current enhanced HOLDDATA again.

Run SMP/E REPORT ERRSYSMODS to identify HIPER/PE exposures. Analyze REPORT ERRSYSMODS output and obtain and process the applicable SYSMODS. Use the IMS Support web site ([techsupport.services.ibm.com/server/390.psp390](http://techsupport.services.ibm.com/server/390.psp390)) or PSP buckets for APAR descriptions. If you find that some PE SYSMODs are already on the system you can:

- Remove the PTF in error if it is not already accepted.
- Leave the PTF in place if the reported PE symptom is not significant.
- Establish operator procedures to restrict access to reported areas of exposure.
- Apply a corrective APAR or PTF, if available.
- Request a FIXTEST from IBM Software Support for the reported problem if the APAR is still open.
- Request a USERMOD code bypass for the reported problem from IBM Software Support.

5. Review the most current IMS PSP buckets for new service information.

6. Test the new maintenance level.

7. Repeat steps 4, 5, and 6 in an ongoing effort to maintain current maintenance levels.

---

## Installing IMS service on a single system

IMS service can be installed in several ways, including SMP/E methods.

The SMP/E methods to install service include:

- RECEIVE/APPLY/ACCEPT (standard sequence)
- ACCEPT without APPLY (pregeneration mode)
- ACCEPT before APPLY (SYSDEF-sensitive service)

### Important:

- Do not ACCEPT APARs or USERMODs.
- Before the installing service, ensure that you have read the latest HOLDDATA information.

If you have any questions about these processes, contact IBM Software Support before you begin.

### Related tasks:

“Installing IMS service in an IMSplex” on page 200

## RECEIVE/APPLY/ACCEPT (standard sequence)

This SMP/E method is the standard and recommended method for processing service.

1. Back up the IMS environment.
  - a. Back up the SMP/E data sets (such as Zones, SMPMTS, and SMPPTS).
  - b. Back up IMS product data sets (such as SDFSRESL and ADFSLOAD).
2. Obtain the desired service.

3. Read the documentation accompanying the package:
  - ESO tape documentation
  - CBPDO Memo to Users Extensions
  - Preventative Service Planning (PSP)
4. Run the SMP/E RECEIVE command with both the parameters SYSMODS and HOLDDATA set.
5. Run the SMP/E APPLY CHECK GROUPEXTEND command.  
 SMP/E Messages GIM43401 and GIM44402 can be received for modules not included in the target system during the APPLY CHECK and APPLY process. You can ignore these messages if they refer to a part that pertains to an IMS function or feature that you are not going to use. Programming exceptions (PEs) need to be resolved to ensure that service is processed to the desired level. If needed, contact IBM Software Support for assistance.  
**Attention:** IMS service frequently includes in-line JCLIN information. For this type of service, SMP/E does not recommend the re-APPLY of service using the REDO parameter. If REDO is used for this type of service (without NOJCLIN), SMP/E RESTORE processing might not work properly.
6. Research the APPLY CHECK reports, making changes as needed.
7. Run the SMP/E APPLY GROUPEXTEND command.
8. Test the corrective service.  
 If an IMS system definition is done with service in APPLY only status, that service might be partially or completely regressed.  
  
**Recommendation:** For all SYSMODs in APPLY only status, issue the following SMP/E command after every IMS system definition, where *xxxx,xxxx* is a list of all SYSMODs in APPLY only status (separated by commas or spaces):  
 APPLY S(*xxxx,xxxx*) REDO NOJCLIN BYPASS (...)
9. Run the SMP/E ACCEPT CHECK GROUPEXTEND command.
10. Research the ACCEPT CHECK reports.
11. Run the SMP/E ACCEPT GROUPEXTEND command.
12. Download a new copy of Enhanced Holddata from the RECEIVE command.
13. Run the SMP/E REPORT ERRSYSMODS command.
14. Analyze the output from the SMP/E REPORT ERRSYSMODS command and process additional services as appropriate.

## ACCEPT without APPLY (pregeneration mode)

This SMP/E method is the pregeneration method for processing service.

**Important:** This information is accurate as of its printing. For the most current and more detailed information, see Information APAR II13024.

This procedure requires that ACCJCLIN was set in the distribution zone when the FMIDs were ACCEPTed.

1. Back up the IMS environment.
  - a. Back up the SMP/E data sets (such as Zones, SMPMTS, and SMPPTS).
  - b. Back up IMS product data sets (such as SDFSRESL and ADFSLOAD).
2. Obtain the desired service.
3. Read the documentation accompanying the package:
  - ESO tape documentation

- CBPDO Memo to Users Extensions
- Preventative Service Planning (PSP)

4. Run the SMP/E RECEIVE command with both the parameters SYSMODS and HOLDDATA set.

5. Run either the SMP/E RESTORE or ACCEPT commands for all outstanding service for all products present in the IMS distribution and IMS target zones. Use the following sample SMP/E control statements to identify this outstanding service (SYSMODS that have been applied only):

```
//SMPCNTL DD *
      SET BDY(targetzonename).
      LIST APAR PTF USERMOD NOACCEPT NOSUP.
```

6. Unload the target zone DDDEFs using the SMP/E UNLOAD command. Use the following sample SMP/E control statements and JCL to complete this task:

```
//SMPCNTL DD *
      SET BDY(targetzonename).
      UNLOAD DDDEF.
//SMPPUNCH DD DSN=IMS.SMPUNLD,DISP=(,CATLG),
// UNIT=SYSDA,SPACE=(TRK,(5,1),RLSE),
// DCB=(RECFM=FB,BLKSIZE=16000,LRECL=80)
```

7. Run the SMP/E LIST command on the target zone to determine the name of the OPTIONS entry. This OPTIONS entry will be used in Step 10. Use the following sample SMP/E control statements to complete this task:

```
//SMPCNTL DD *
      SET BDY(targetzonename).
      LIST TARGETZONE.
```

8. Scratch and reallocate the following data sets:

- SMPMTS
- SMPSTS
- SMPSCDS
- SMPLTS

**Note:** The SMP/E CLEANUP command can be used instead of scratching and reallocating the SMPPTS, SMPSTS and SMPSCDS data sets. The SMP/E CLEANUP command cannot be used for the SMPLTS data set. You need to scratch and reallocate this data set.

Use the following sample SMP/E control statements to complete this task:

```
//SMPCNTL DD *
      SET BDY(targetzonename).
      CLEANUP.
```

9. Delete the SMP/E target zone.

a. Run the SMP/E ZONEDELETE command for the Target zone. Use the following sample SMP/E control statements to complete this task:

```
//SMPCNTL DD *
      SET BDY(targetzonename).
      ZDEL TZONE(targetzonename).
```

b. If no other SMP/E zones are in the target CSI (the VSAM cluster), run the IDCAMS DELETE and DEFINE commands on the target CSI to improve performance.

**Attention:** If multiple zones are contained in the same CSI as the target zone, do not delete and redefine the cluster because you will also lose the information for those zones.

10. Re-initialize the new Target zone.

- a. Run the IDCAMS REPRO command to copy SYS1.MACLIB(GIMZPOOL) into the new CSI.

**Attention:** If you did not delete and redefine the target CSI as described in Step 9 on page 197, do not copy GIMZPOOL into the new target zone.

- b. Rebuild the relationship between the old DLIB zone and the new Target zone. Use the following sample SMP/E control statements to complete this task:

```
//SMPCNTL DD *
    SET BDY(GLOBAL).
    UCLIN.
    ADD GZONE ZONEINDEX(
        (targetzonename,target.zone.cluster.name,TARGET)
    ).
    ENDUCL.

    SET BDY(targetzonename).
    UCLIN.
    ADD TARGETZONE(targetzonename)
    SREL(P115)
    RELATED(dlibzonename)
    OPTIONS(xxxxxx).
    ENDUCL.
```

**Note:** Be sure that this new target points to the correct OPTIONS entry. The correct OPTIONS entry can be determined from the output created in step 7 on page 197.

- c. Run UCLIN to add the DDDEFs back to the target zone. This step uses the data set created in step 6 on page 197 as input. Use the following SMP/E control statements and JCL to complete this task:

```
//SMPCNTL DD *
    SET BDY(targetzonename).
    //      DD DSN=IMS.SMPUNLD,DISP=SHR
```

**Note:** Return code 4 is expected in this step because DDDEFs are being added instead of being replaced.

**Attention:** Before processing SMP/E in step 11, RECEIVE the current Enhanced HOLDDATA. Doing so enables you to resolve PEs during SMP/E processing. You can get the most recent HOLDDATA at <http://service.boulder.ibm.com/390holddata.html>.

11. Run the SMP/E ACCEPT GROUPEXTEND BYPASS(APPLYCHECK) commands for the PTFs to be processed. Use the following sample SMP/E control statements to complete this task:

```
//SMPCNTL DD *
    SET BDY(dlibzonename).
    ACCEPT GROUPEXTEND
    BYPASS(APPLYCHECK
        HOLDCLASS(ERREL,UCLREL)
        HOLDSYSTEM
    )
    SOURCEID (SMCREC,RSU08*,RSU090*,RSU0910,etc)
    PTFS.
```

**Attention:** After SMP/E processing is complete, using the current enhanced HOLDDATA, run the SMP/E REPORT ERRSYSMODS command to identify missing HIPERs and PE exposures.

12. Run the SMP/E ZONEMERGE command specifying CONTENT to merge the distribution zone to the new target zone. Use the following sample SMP/E control statements to complete this task:

```
//SMPCNTL DD *
      SET BDY(targetzonename).
      ZONEMERGE(dlibzonename)
      INTO(targetzonename)
      CONTENT.
```

13. Run the SMP/E GENERATE command to create the JCL necessary to rebuild the target libraries.

**Tip:** This step requires that ACCJCLIN was set in the distribution zone before the IMS FMIDs were ACCEPTed.

Use the following sample SMP/E control statements and additional JCL to complete this task, where the data set for DD name CNTL must have a member named J, which contains a sample JOB card:

```
//CNTL      DD DSN=yourpds,DISP=SHR
//SMPPUNCH DD DSN=IMS.GENERATE,
//           DISP=(,CATLG),UNIT=SYSDA,
//           SPACE=(CYL,(25,5),RLSE),
//           DCB=(RECFM=FB,LRECL=80,BLKSIZE=16000)
      SET BDY(targetzonename).
      GENERATE JOBCARD(CNTL,J) REPLACE.
```

14. Run the JCL that was created in step 13.

**Note:** The SMPLTS job will complete with a return code of 4 because of unresolved external references (IEW2454W). All other jobs should complete with a return code of 0.

15. Run an IMS ALL type of system definition (STAGE 1 and STAGE 2).

**Requirement:** If you are running a large IMS system (LGEN), you need to allocate the IMS.LGENIN and IMS.LGENOUT data sets before running the IMS SYSDEF preprocessor. The preprocessor performs Stage 1 for a large system definition.

**Attention:** Ensure that Stage 2 processing is complete before performing step 16.

16. Run SMP/E JCLIN, pointing to the STAGE 2 JCL as input.

**Requirement:** If you are running a large IMS system, the output from the DFSIVG10 program must be used as input to JCLIN. For assistance with DFSIVG10, refer to IVP job IV\_C401J.

17. Run the SMP/E APPLY command for any IMS service that was not accepted. This service was identified in step 5 on page 197.
18. Run the SMP/E APPLY command for any service for other products that was not accepted. This service was identified in step 5 on page 197.
19. Download a new copy of Enhanced Holddata from the RECEIVE command.
20. Run the SMP/E REPORT ERRSYSMODS command.
21. Analyze the output from the SMP/E REPORT ERRSYSMODS command and process additional services as appropriate.
22. Test the new system.

## ACCEPT before APPLY (for service that requires a system definition)

This method is a variation of pregeneration mode that can be useful when you have many products sharing the same SMP/E zones and you need to install a PTF that would normally require an ACCEPT BYPASS(APPLYCHECK) sequence (typically a PTF that affects system definition).

This method avoids disturbing other products that have outstanding service (service that has been APPLIED but not ACCEPTed).

1. Back up the IMS environment.
  - a. Back up the SMP/E data sets (such as Zones, SMPMTS, and SMPPTS).
  - b. Back up IMS product data sets (such as SDFSRESL and ADFSLOAD).
2. Obtain the desired service.
3. Read the documentation accompanying the package:
  - ESO tape documentation
  - Preventative Service Planning (PSP)
4. Run the SMP/E RECEIVE command with both the parameters SYSMODS and HOLDDATA set.
5. Run the SMP/E ACCEPT or RESTORE commands on outstanding APPLY service for all products sharing the SMP/E zones with IMS.
6. Run the SMP/E ACCEPT CHECK GROUPEXTEND BYPASS(APPLYCHECK) command.
7. Research the ACCEPT CHECK reports, making changes as necessary.
8. Run the SMP/E ACCEPT GROUPEXTEND BYPASS (APPLYCHECK) command.
9. Run an IMS ALL type of system definition STAGE 1 and STAGE 2.
10. Run SMP/E JCLIN pointing to the STAGE 2 JCL as input.
11. Run SMP/E APPLY CHECK GROUPEXTEND.  
**Attention:** Do not use the REDO parameter.
12. Research the APPLY CHECK reports, making changes as necessary.
13. Run SMP/E APPLY GROUPEXTEND.  
**Attention:** Do not use the REDO parameter.
14. Download a new copy of Enhanced Holddata from the RECEIVE command.
15. Run the SMP/E REPORT ERRSYSMODS command.
16. Analyze the output from the SMP/E REPORT ERRSYSMODS command and process additional services as appropriate.
17. Test the new system.

---

## Installing IMS service in an IMSplex

Installing IMS service in an IMSplex is similar to installing service on a single system.

The following considerations apply when planning to install service into an IMSplex:

- Multiple Resource Managers (RMs) and Operations Managers (OMs) can run simultaneously in an IMSplex.  
If all your IMS systems on one LPAR, one OM must be running at all times to provide OM services.
- Only one Structured Call Interface (SCI) can run at any given time in an IMSplex on a single logical partition (LPAR).
- All Common Queue Server (CQS) clients connected to a CQS address space must be stopped before shutting down that CQS.

### Recommendations:



- For continuous availability, do not install service to the entire IMSplex at the same time.
- If you are running multiple IMS systems on one logical partition (LPAR), treat each IMS system within the IMSplex as a separate system, bringing each down individually, implementing service, testing it, and bringing it back online. After each system has successfully been brought back online, rotate to the next system, repeating the process.  
Apply service to less complex IMS systems before applying service to more complex IMS systems.
- If you are running multiple LPARs, install service on one LPAR at a time.  
Apply service to less complex LPARs before applying service to more complex LPARs.
- If your IMSplex includes IMS Connect, ensure that all appropriate IMS systems are available before starting IMS Connect. Otherwise, IMS Connect issues a “datastore unavailable” message.

**Related tasks:**

“Installing IMS service on a single system” on page 195

---

## Common installation and maintenance issues

You can have a more stable IMS environment by being aware of some of the common installation and maintenance issues that are presented in these topics and taking appropriate action.

### Preventing regression of SYSMODs in APPLY-only status by an IMS system definition

If an IMS system definition is performed when maintenance is in APPLY status, the maintenance might be regressed.

There are two methods of preventing SYSMODs in this status from being regressed by an IMS system definition:

- Before the system definition, ACCEPT all PTFs in Apply-only status.
- After the system definition, reprocess the Apply-only SYSMODs.

Use the following SMP/E statements to list the SYSMODs in Apply-only status:

```
SET BOUNDARY (targetzone).
LIST APAR PTF USERMOD NOACCEPT NOSUP.
```

**Attention:** Only use the NOJCLIN parameter when processing REDO. Otherwise, you might not be able to RESTORE the service.

Use the following SMP/E statements to reprocess SYSMODs in Apply-only status

```
APPLY REDO NOJCLIN SELECT(
    xxxxxx
    xxxxxx
    xxxxxx
) .
```

xxxxxx explicitly lists each SYSMOD in Apply-only status. NOJCLIN should only be specified for REDO processing. Some SYSMODs in Apply-only status might have System HOLDS for system definition that require special handling. Follow the instructions in the HOLD statements for those SYSMODs.



## Generating JCL to build non-system definition target libraries

Some elements of IMS are not included in the IMS system definition process. These elements are identified to SMP/E and built during APPLY processing for their FMIDs.

The SMP/E GENERATE command can be used to create JCL that can be used to rebuild these components in their target libraries. SMP/E GENERATE can also be used to create JCL for other products in the IMS distribution zone, such as IRLM. SMP/E GENERATE processing is dependent on the SMP/E parameter ACCJCLIN being set in the distribution zone when the FMID is ACCEPTed.

## Applying maintenance for the IVP dialog

Service affecting the IVP dialog process can require that special processing be performed.

SMP/E HOLDDATA identifies the required actions, if there are any that need to be performed.

The following actions might need to be performed, as identified in HOLDDATA:

- Table Merge  
Table Merge is necessary if rows have been added, changed, or deleted in one of the master tables. Table merge causes the changes to be propagated to the user tables in INSTALIB. Default values for variables are not updated for variables that have been changed by dialog processing.
- Variable Gathering  
You can modify the default values for new and changed variables.
- File Tailoring  
You can rerun File Tailoring to add INSTALIB members for new JOBS or TASKs or to update INSTALIB members with new or changed variable values.
- Execution  
You can run or rerun portions of the IVP processes.

## Upgrading z/OS

Before you upgrade the z/OS system on which IMS is running, consider any requirements for the z/OS interface and VTAM interfaces.

Ensure that you perform the following when upgrading the z/OS system on which IMS is running:

1. Review the z/OS considerations. See Chapter 5, “z/OS interface considerations,” on page 125 for more information.
2. Review the VTAM considerations. See Chapter 6, “VTAM interface considerations,” on page 139 for more information.

**Note:** The IVP D series of samples contains examples of all of the z/OS and VTAM interfaces, except for the Channel-to-Channel (CTC) Channel-End Appendage. See *IMS Version 11 Installation* for more information.

## Ensuring proper SYSLIB concatenation

The order in which your macro libraries are concatenated is critical. Otherwise, unpredictable results might occur during assembler processing. Ensure that your libraries are concatenated in the correct order.

### **SMP/E Apply:**

1. IMS.OPTIONS
2. SMPMTS
3. IMS.SDFSMACT
4. MVS Macro Libraries

### **IMS.OPTIONS**

Built during STAGE 2 of system definition and is customized by the specified parameters.

### **SMPMTS**

A target library for macros that exist only in a distribution library. This data set enables the current version of the macros to be used for assemblies during APPLY processing.

### **IMS.SDFSMACT**

Target library for all IMS macros.

### **MVS Macro Libraries**

Consist of the appropriate combination of SYS1.MACLIB (AMACLIB), SYS1.MODGEN (AMODGEN), SYS1.TSOMAC (ATSOMAC), and ASM.SASMMAC2. ASM.SASMMAC2 contains concept 14 macros and comes with the High Level Assembler in the HLASM Toolkit.

### **SMP/E Accept:**

1. IMS.OPTIONS
2. IMS.ADFSMACT
3. MVS Macro Libraries

**Note:** The noticeable difference from the SMP/E APPLY process is the absence of SMPMTS, and SMP/E pointing to distribution libraries rather than target libraries. SMPMTS contains versions of macros that have not been accepted.

### **IMS System Definition Stage 1:**

1. IMS.ADFSMACT

### **IMS System Definition Stage 2:**

1. IMS.OPTIONS
2. IMS.ADFSMACT
3. MVS Macro Libraries

## **Interpreting binder return codes properly**

Some binder return codes can be safely ignored while others cannot.

The following table lists the acceptable return codes from the various binder processes:

*Table 12. Acceptable return codes from the binder*

| Type of SMP/E processing | Return code | How to interpret                                                                  |
|--------------------------|-------------|-----------------------------------------------------------------------------------|
| APPLY                    | 0           | Do not ignore unresolved external references. The exception is binds into SMPLTS. |
| ACCEPT                   | 4           | You can safely ignore unresolved external references.                             |

Table 12. Acceptable return codes from the binder (continued)

| Type of SMP/E processing  | Return code | How to interpret                              |
|---------------------------|-------------|-----------------------------------------------|
| System Definition STAGE 2 | 0           | Do not ignore unresolved external references. |

**Recommendation:** Point to a different utility entry in SMP/E for APPLY and ACCEPT processing.

## Assembling and binding a sample exit routine using SMP/E

The following example demonstrates a technique that you can use to have SMP/E assemble and bind one of the sample exit routines.

```

++ USERMOD (XYZUMOD) .
++ VER (P115)
   FMID(HMK1100) .
++ JCLIN.
//INJCLIN JOB . . .
//LKED EXEC PGM=IEWL,
//      PARM=('SIZE=(880K,64K)',RENT,REFR,NCAL,LET,XREF,LIST)
//ADFSLOAD DD DSN=IMS.ADFSLOAD,DISP=SHR
//SYSPUNCH DD DSN=IMS.OBJDSET,DISP=SHR
//SYSUT1 DD UNIT=(SYSDA,SEP=(SYSLMOD,SYSLIN)),SPACE=(1024,(200,20))
//SYSPRINT DD SYSOUT=A
//SYSLMOD DD DSN=IMS.SDFSRESL,DISP=SHR
//SYSLIN DD *
   INCLUDE ADFSLOAD(DFSCSI00)
   INCLUDE SYSPUNCH(DFSGMSG0)
   ENTRY DFSGMSG0
   NAME DFSGMSG0(R)
++ SRC (DFSGMSG0) SYSLIB(SDFSSMPL) DISTLIB(ADFSSMPL) .
DFSGMSG0 TITLE 'DFSGMSG0 -- GREETING MESSAGES user exit routine routine'
.....
.....
.....

```

## Migrating to a new version of IMS

When migrating to a new version of IMS, there are certain tasks that should be performed regardless of which version you are migrating from.

When migrating to a new version of IMS, ensure that you perform the following tasks:

1. Review the *Release Planning Guide* for the version that you are migrating to. In particular, review the migration and coexistence information.
2. If you are skipping a version, review the *Release Planning Guide* for those versions you are skipping. In particular, review the migration and coexistence information.
3. Review the PSP bucket for the version that you are migrating to.
4. If you are skipping a version, review the PSP bucket for those versions you are skipping.

---

## Chapter 14. Planning for shared queues

This topic gives an overview of the concepts you should understand to use shared queues and then outlines the planning and administrative tasks associated with shared queues.

To manage message queues, your installation can use the IMS queue manager with either message queue data sets or shared queues.

**Restriction:** If you use shared full-function message queues on a Fast Path-eligible system, you must also define shared queues for EMH messages.

Before operating in a shared-queues environment, read the information in this topic, which explains how operating in a shared-queues environment differs from operating in a non-shared-queues environment.

---

### The role of CQS in a shared-queues environment

The Common Queue Server (CQS) is a generalized server that manages objects on a coupling facility list structure, such as a queue structure or a resource structure. CQS receives, maintains, and distributes data objects from shared queues on behalf of multiple clients.

Each client communicates with a CQS to access the shared queues. IMS is one example of a CQS client that uses CQS to manage both its shared queues and shared resources.

CQS uses the z/OS coupling facility as a repository for data objects. Storage in a coupling facility is divided into distinct objects called structures. Authorized programs use structures to implement data sharing and high-speed serialization. The coupling facility stores and arranges the data according to list structures. Queue structures contain collections of data objects that share the same name, known as queues.

### The IMS role in a shared-queues environment

In a shared-queues environment, IMS acts as a front end for terminals, as a back end for processing, or as both simultaneously.

#### IMS as a front end

As a front end, an IMS manages terminal resources and originates communications traffic, such as commands, transactions, and message switches. The message is typically the result of data that is entered at a terminal. A front-end IMS can also play the role of back end, if database, program, and transaction resources are defined.

#### IMS as a back end

As a back end, an IMS manages transactions, database resources, and schedules application programs. A back-end IMS can also play the role of front end, if terminal resources are defined.

## IMS as both a front end and a back end

As both a front end and a back end, IMS can process messages that are entered from one IMS terminal either on that IMS or on another IMS within the IMSplex.

Three processing environments exist when IMS acts as a back-end processor: local, remote MSC, and remote shared queues.

### Definitions:

- *Local processing* occurs when the IMS dependent region that processes the message is on the same IMS as the IMS control region that received the message.
- *Remote MSC processing* occurs when the IMS dependent region that processes the message is on a different IMS than the IMS control region that received the message, and the message is sent to the remote IMS using MSC.
- *Remote shared-queues processing* occurs when the IMS dependent region that processes the message is on a different IMS than the IMS control region that received the message. In this case, however, the shared-queues facility contains the message that the remote IMS processes.

## How IMS registers interest in queues

An IMS system registers and deregisters its interest in shared queues according to the type of work the IMS system can process. Consequently, the CQS for this IMS can notify the IMS when work is present on a queue.

The types of work include:

- Transactions
- LTERMs
- MSC resources (such as remote LTERMs, remote transactions, and MSNAMEs)

In general, IMS registers its interest as the resources are added (or started), and IMS deregisters its interest as the resources are deleted (or stopped).

### Queue types

Various types of queues are managed in a shared-queues environment. Each queue type is used for a different type of work. An IMS registers interest in only those queue types that it can process, based on the types of work you define for it.

The queue types and the work that IMS processes on them are shown in the following table.

*Table 13. Queue types maintained in a shared-queues environment*

| Queue type                | Description                                                                                                                                    |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Transaction-ready queue   | Contains first qbuffer of messages that are destined for transactions                                                                          |
| Transaction-staging queue | An internal queue that the IMS Queue Manager uses for holding those parts of messages that exceed a single qbuffer                             |
| Transaction-suspend queue | Contains first qbuffer of messages destined for suspended transactions                                                                         |
| Transaction-serial queue  | Contains first qbuffer of messages that are destined for serial transactions (transactions defined to IMS on the TRANSACT macro as SERIAL=YES) |
| LTERM-ready queue         | Contains first qbuffer of messages that are destined for LTERMs or MSNAMEs                                                                     |

Table 13. Queue types maintained in a shared-queues environment (continued)

| Queue type          | Description                                                                                                        |
|---------------------|--------------------------------------------------------------------------------------------------------------------|
| LTERM-staging queue | An internal queue that the IMS Queue Manager uses for holding those parts of messages that exceed a single qbuffer |
| APPC-ready queue    | Contains first qbuffer of messages that are destined for APPC devices                                              |
| REMOTE-ready queue  | Contains first qbuffer of messages that are destined for remote transactions or remote LTERMs                      |
| OTMA-ready queue    | Contains first qbuffer of messages that are destined for OTMA devices                                              |

## Registering and deregistering interest in transactions

IMS registers or deregisters interest in transaction queues when certain actions are taken.

- IMS registers interest in transaction queues when any of the following events occurs:
  - IMS is initialized.
  - Transactions are added by online change after the /MODIFY COMMIT command is entered or, with global online change enabled, after the INITIATE OLC PHASE(COMMIT) command is entered.
  - An operator enters the /START TRAN command.
  - An operator enters the UPD TRAN START(SCHD) command.
  - IMS reconnects to CQS after CQS has terminated and been restarted.
- IMS deregisters interest in transaction queues when any of the following events occurs:
  - IMS shuts down.
  - Transactions are deleted by online change after the /MODIFY COMMIT command or, with global online change enabled, after the INITIATE OLC PHASE(COMMIT) command.
  - A /STOP TRAN command is entered.
  - A /PSTOP TRAN command is entered.
  - An UPD TRAN STOP(SCHD) command is entered.
- For IMS Fast Path:
  - IMS registers interest in a program queue during IFP region initialization for the name of the program that the region processes.
  - IMS deregisters interest in a program queue when the IFP region is terminated.

## Registering and deregistering interest in LTERMs

IMS registers and deregisters interest in LTERM queues when certain actions are taken.

- IMS registers interest in LTERM queues when any of the following events occurs:
  - A user logs on to a static terminal.
  - A user signs on to a dynamic terminal.
  - An LTERM is assigned to an active user or node.
- IMS deregisters interest in LTERM queues when any of the following events occurs:
  - A user logs off from a static terminal.

- A user signs off from a dynamic terminal.
- An LTERM is assigned to an inactive user or node.

### Registering interest in MSC resources

When an MSC logical link is started, IMS registers interest in the MSNAME that is defined with it.

All remote transactions that are defined using the same SYSID pair are also registered.

When the logical link is stopped, IMS deregisters interest in the MSNAME and its associated remote transactions.

## Message flow in a shared-queues environment

In a shared-queues environment, when an IMS retrieves a message from a shared queue, the message is locked on that shared queue until the IMS unlocks or deletes the message. If IMS does not unlock or delete a message from the shared queue (for example, because of an IMS abend), the message remains locked.

Locked messages are not available to other IMS systems for processing. If you cold start the IMS after it locks messages on a shared queue, those messages remain locked, and they are moved to the cold queue.

If RM is defined with a resource structure, IMS releases any locked messages destined for an LTERM on the failing IMS. Messages become available to a user signing on to another IMS in the IMSplex.

**Definition:** The *cold queue* is a type of queue on which CQS places locked messages that it finds on its private queues after a client cold starts. Messages remain on the cold queue until they are unlocked or deleted.

## Units of work (UOW) tracking

In a non-shared-queues environment, a unit of work is tracked by a recovery token and by a unit of work (UOW). In a shared-queues environment, a unit of work is tracked using only a UOW.

The UOW consists of the following fields:

- Originating-system message ID: Message ID assigned by the IMS that originates the message.
  - Originating IMSID
  - Time-stamp token
- Processing-system message ID: Message ID assigned by the IMS processing the message.
  - Processing IMSID
  - Time-stamp token

Because the UOW has IDs for both the system that originates the message and the system (if any) that processes the message, all messages that are associated with an original message can be tied together by the UOW (specifically, the originating-system message ID in the UOW).

## Terminal autologon

Application programs can run on any IMS back-end system; therefore, output can be generated for the same user or terminal on any of these systems.

In order to deliver output to all waiting users, an autologon terminal is often switched between IMS systems in the IMSplex.

## Message and program-to-program switches

Message switches and non-express PCB program-to-program switches are queued globally in a shared queue environment. Program-to-program switches to express PCBs are eligible to be considered for local-first optimization. If the express-PCB message is not selected for local-first processing, then it is queued globally.

If the destination of a message switch is not defined on the local IMS, the Destination Creation exit routine (DFSINSX0) is called to identify the destination.

While ETO is active, if the exit routine indicates that the destination is an LTERM, the exit routine creates a dynamic user structure for the LTERM and then queues the message.

The DFSINSX0 exit can also indicate that the destination is a transaction. IMS will create the transaction based on the environment and options provided by the exit:

- In a shared-queues environment without DRD, IMS will create a queuing-only transaction and queue the message globally.
- In a shared-queues environment with DRD enabled, IMS will create a transaction either for queueing or for scheduling, as requested by the exit.

## Dynamic control blocks

To enable a user to enter transactions on IMS systems that do not have the transaction defined, a user can supply information in the Destination Creation exit routine (DFSINSX0) to cause IMS to build a dynamic scheduler message block (SMB).

Input and output messages are enqueued to the shared queues, rather than to the control blocks.

## IMS Queue Manager

The IMS Queue Manager manages the IMS message queues. The Queue Manager data sets, SHMSG and LMSG, are not used to back up messages in a shared-queues environment.

However, in order to reduce necessary changes during IMS system definition, the MSGQUEUE macro is still required; the data sets are defined, but they are never allocated or opened.

The Queue Manager uses a number of real storage queue buffers to store the messages prior to placing them in the shared queues on the coupling facility. However, if the message needs to remain in the local IMS rather than being placed on the shared queues, the message is placed on a local queue. The local queues are also managed by the Queue Manager.

Messages that are placed on the shared queues are recoverable. Because the message queue data sets are not used, a message on a local queue is only



recoverable if the log record for the message was written to the IMS log after the oldest of the last two checkpoints was taken.

In an XRF environment, the local queue manager data sets (QBLKSL, SHMSGL, and LGMSGL) are used. While the XRF alternate subsystem is tracking, local messages are placed in the local message queue data sets (messages for the active subsystem are on the shared queues). When the XRF alternate subsystem takes over the active subsystem, any messages in the local message queue data sets are merged with those on the shared queues.

## MTO messages

Messages for the primary master terminal operator (MTO) are kept local and are not recovered across an IMS restart. Messages for the secondary MTO are placed on the shared queues using the LOCAL=YES option.

These secondary MTO messages are recovered across an IMS restart and not are retrieved by the CQSREAD.

## Serial application processing

In an IMSplex environment that does not use shared queues, if you want to define an application program as a serial program (that is, it cannot be scheduled simultaneously into message or batch message regions), you use the SCHDTYP= parameter on the APPLCTN macro. Application programs defined as serial cannot be scheduled to run simultaneously in more than one message or batch message region. Parallel application programs (SCHDTYP=PARALLEL), however, can be scheduled to run simultaneously.

In an IMSplex environment with shared queues, you can also define an application program as serial. In addition to defining the application program as serial (SCHDTYP=SERIAL), you must:

- Define an active instance of Resource Manager.
- Update the CQS global structure definition IMS PROCLIB member data set (CQSSGxxx) to define a resource structure using the RSRCSTRUCTURE parameter.
- Identify the resource structure using the RSRCSTRUCTURE parameter of the CSL RM initialization parameters PROCLIB member data.

When shared queues is enabled, and RM is active with a defined resource structure, IMS serializes the scheduling of applications defined as serial in an IMSplex.

### **Related reading:**

- For information on defining and tailoring RM, see *IMS Version 11 System Definition*.
- For information on the CQSSGxxx IMS PROCLIB member data set, see *IMS Version 11 System Definition*.

## Serial transaction processing

Messages for serial transactions are only processed by the IMS that places the messages on the queue. IMS keeps a transaction-serial queue for serial transaction messages in the coupling facility MSGQ structure.

AOI programs must define their transactions as serial (specified as SERIAL=YES on the TRANSACT macro).

## Conversational transactions in a shared-queues environment

In a shared-queues environment, conversation status is local to the IMS that initiates the conversation.

When an IMS initiates a conversation, the following events occur:

1. IMS places the conversational transaction on the shared queues.
2. A locally defined SMB (or RSMB) or the Output Creation User exit routine (DFSINSX0) identifies the transaction as conversational.
3. Any IMS that registers interest in the transaction can process it. Also, any IMS can process any step of the conversation.
4. One IMS reads the message and schedules an application program to process the message.
5. The application program processes the message, saves data in the SPA, and responds to the initiating terminal.

The /EXIT, /HOLD, and /RELEASE commands apply only to the IMS that initiates the conversation.

---

## Configuring a shared-queues environment

You can configure a shared-queue environment by using either a cloned configuration or a partitioned configuration.

Figure 27 on page 212 and Figure 28 on page 214 illustrate two methods of configuring a shared-queues environment.

### Cloned configuration

A *cloned configuration* involves creating identical IMS systems, with almost identical system definitions. Some differences in the definitions are required, such as MSC links. You can override the MTO definitions generated during system definition in IMS PROCLIB member data set DFSDCxxx.

**Related reading:** For more information on overriding the MTO definitions, see *IMS Version 11 System Definition*.

The LTERMs, databases, and transactions for all the IMS systems are defined identically. In this configuration, any IMS is able to process a particular transaction. Cloned configurations are useful for automatic workload balancing.

**Recommendation:** To take full advantage of a cloned configuration, ensure that the databases are shared across the IMSplex.




Figure 27. Cloned configuration in a shared-queues environment

## Partitioned configuration

A *partitioned configuration* separates the VTAM network responsibilities from the database work. Each resource is defined to only one IMS.

You can define an ES/9000 9021 to manage the VTAM network as an IMS front end, while multiple smaller 9672 CMOS processors manage the database work.

Partitioned configurations are useful as replacements for, or additions to, MSC networks.

In the partitioned configuration in the following figure, IMS A acts as the front end, and has LTERM 1, LTERM 2, and MTO 1 defined. IMS B, IMS C, and IMS D act as back ends, each with separate transactions and databases defined. Transactions in this configuration can run only on the IMS systems on which they are defined.




Figure 28. Partitioned configuration in a shared-queues environment

## Enabling shared queues

You enable shared queues by defining them in the z/OS coupling facility resource management (CFRM) policy.

This topic discusses what you must include in the CFRM policy and the other definitions that must conform to the definitions you make in the CFRM policy.

## Defining the CFRM policy

To enable shared queues, you must define a CFRM policy that contains structure names and attributes for all message queues. If you use Fast Path and share EMH queues, you must also include structure names and attributes for EMHQ structures.

Be sure that the structure names you define in the CFRM policy for message queues and any optional EMH queues are also defined in the following places:

- The shared-queues IMS.PROCLIB member data set (DFSSQxxx) parameters MSGQ= (and EMHQ= if you are sharing EMH queues)
- The CQS local structure definition PROCLIB member data set (CQSSLxxx) parameter STRNAME=
- The CQS global structure definition PROCLIB member data set (CQSSGxxx) parameters STRNAME= and OVFLWSTR=

**Recommendation:** If possible, place high-use structures on separate coupling facilities. For example, place your IMS data-sharing structures on a different coupling facility from your shared-queues structures. Place your VSAM and DB2 for z/OS structures on yet another coupling facility. You can place low-use structures (such as those for RACF) on any coupling facility where space is available.

**Related reading:** For more information on definitions that are required to use CQS, see *IMS Version 11 System Definition*.

The following example shows a sample CFRM policy containing definitions of shared queues and a resource structure.

```
//CFRMPLCY JOB MSGCLASS=A,REGION=2000K,CLASS=K
// MSGLEVEL=(1,1)
//*****
/* This JCL is used for configuration. INITSIZE is *
/* used for the primary MSGQ and EMHQ structures. *
//*****
/* 2 CF *
//*****
//POLICY EXEC PGM=IXCM2APU
//STEPLIB DD DSN=SYS1.MIGLIB,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSIN DD *

DATA TYPE(CFRM)
DEFINE POLICY
    NAME(CONFIG01)
    REPLACE(YES)

    CF NAME (CF01)
    TYPE(nnnnnn)
    MFG(aa)
    PLANT(nn)
    SEQUENCE(nnnnnnnnnnnn)
    PARTITION(n)
    CPCID(nn)

    CF NAME (CF02)
    TYPE(nnnnnn)
    MFG(aa)
```

```

    PLANT(nn)
    SEQUENCE(nnnnnnnnnnnn)
    PARTITION(n)
    CPCID(nn)
    .
    .
    .

STRUCTURE NAME(QMSGIMS01)
    SIZE(16000)
    INITSIZE(8000)
|    MINSIZE(8000)
    PREFLIST(CF01,CF02)
    REBUILDPERCENT(1)
|    ALLOWAUTOALT(YES)
    FULLTHRESHOLD(60)

STRUCTURE NAME(QMSGIMS010FLW)
    SIZE(8000)
|    MINSIZE(8000)
    PREFLIST(CF01,CF02)
    REBUILDPERCENT(1)
|    ALLOWAUTOALT(YES)
    FULLTHRESHOLD(60)

STRUCTURE NAME(QEMHIMS01)
    SIZE(16000)
    INITSIZE(10000)
|    MINSIZE(10000)
    PREFLIST(CF01,CF02)
    REBUILDPERCENT(1)
|    ALLOWAUTOALT(YES)
    FULLTHRESHOLD(60)

STRUCTURE NAME(QEMHIMS010FLW)
    SIZE(8000)
|    MINSIZE(8000)
    PREFLIST(CF01,CF02)
    REBUILDPERCENT(1)
|    ALLOWAUTOALT(YES)
    FULLTHRESHOLD(60)

STRUCTURE NAME(MVSLOGQMSG01)
    SIZE(16000)
    INITSIZE(11000)
    PREFLIST(CF01,CF02)

STRUCTURE NAME(MVSLOGQEMH01)
    SIZE(4000)
    PREFLIST(CF01, CF02)
    REBUILDPERCENT(1)

STRUCTURE NAME(QRSCIMS01)
    SIZE(16000)
    INITSIZE(8000)
|    MINSIZE(8000)
    ALLOWAUTOALT(YES)
    FULLTHRESHOLD(60)
    DUPLEX(ALLOWED)
    PREFLIST(CF01,CF02)
    .
    .
    .

```

*Defining message queue list structures*

You must define the primary list structure, MSGQ, to IMS before using shared queues. You can optionally define the overflow structure.

In the example, STRUCTURE NAME(QMSGIMS01) defines a message queue structure, and STRUCTURE NAME(QMSGIMS010FLW) defines an overflow structure for the QMSGIMS01 message queue.

### *Defining a resource structure*

You must define a resource structure to use RM to share global status. In the example, STRUCTURE NAME(QRSCIMS01) defines a resource structure. For more information about how RM maintains global resource information, see “Information managed by the CSL RM” on page 169.

### *Defining Fast Path message queue list structures*

If you define the MSGQ primary list structure on an IMS that has Fast Path defined, you can optionally define a primary list structure for shared EMH queues. This structure is called the EMHQ structure. You can then optionally define an EMHQ overflow structure.

If you define an EMHQ structure, Fast Path transactions can be processed by IMS systems in the shared-queues group. If you do not define an EMHQ structure, Fast Path transactions, if any, are only processed locally.

In the example, STRUCTURE NAME(QEMHIMS01) defines an EMHQ structure, and STRUCTURE NAME(QEMHIMS010FLW) defines an overflow structure for the QEMHIMS01 EMH queue.

## **Defining a z/OS log stream**

You must define a z/OS log stream. You define a z/OS log stream by updating definitions.

To define a z/OS log stream:

- In the GRSTNLxx PARMLIB member, specify every IMS in the Parallel Sysplex environment as being in the same global resource serialization complex. This allows global serialization of resources in the IMSplex.
- Define log stream structure names and attributes in the CFRM policy for message queues and, if sharing, EMH queues.

**Example:** In the example in “Defining the CFRM policy” on page 215, STRUCTURE NAME(MVSLOGQMSG01) defines a log stream structure name for the message queues. Similarly, STRUCTURE NAME(MVSLOGQEMH01) defines a log stream structure for Fast Path EMH queues.

- Install DFSMS and change SMS DATACLAS, STORCLAS, and MGMTCLAS constructs.
- Format the inventory coupled data set and define log streams and structure names.
- Update the COUPLExx PARMLIB member.
- Define the logger inventory data set in SYS1.PARMLIB(COUPLExx).
- Define log stream and coupling facility structure names to the logger inventory coupled data set.



**Related reading:** For more information on defining the z/OS log stream, see *IMS Version 11 System Definition*.

## Defining CQS parameters

The CQSIPxxx, CQSSLxxx, and CQSSGxxx members of the PROCLIB data set must be updated to enable shared queues.

To enable shared queues, define the following CQS parameters:

- CQS initialization parameters in PROCLIB member data set CQSIPxxx.
- CQS local structure definition parameters in PROCLIB member data set CQSSLxxx (for local structure definition). These definitions apply only to a single CQS.
- CQS global structure definition parameters in PROCLIB member data set CQSSGxxx (for global structure definition).

You must define one of each of these parameters for each structure pair. All of the CQs that share queues must have the same values specified in the CQSSGxxx PROCLIB member data set. If they differ, only the first CQS connects to the structure, and the others fail.

- Execution parameters (optional).

If you do not specify ARMRST=, CQSGROUP=, SSN=, STRDEFG=, or STRDEFL=, CQS retrieves the values that are specified in the CQSIPxxx PROCLIB member data set.

**Related reading:** For more information about the CQS parameters, see *IMS Version 11 System Definition*.

## IMS parameters for shared queues

To enable shared queues, you must define IMS control region parameters and IMS startup parameters.

Define the following IMS control region execution parameters:

### **LGMSGSZ=**

The record length of the long message.

### **QBUF=**

The upper expansion limit of the queue buffer pool.

### **QBUFHITH=**

A one- to three-digit number with a value in the range of 1 to 100 that represents the high threshold percentage at which the message queue buffer is dynamically expanded. The default percentage is 80%.

### **QBUFLWTH=**

A one- to three-digit number with a value in the range of 1 to 100 that represents the low threshold percentage at which the message queue buffer is compressed. The default percentage is 50%.

### **QBUFMAX=**

The maximum number of message queue buffers that are allowed in the queue buffer pool.

### **QBUFPCTX=**

A one- to three-digit number with a value in the range of 1 to 100 that

represents the percentage of the originally allocated message queue buffers that are dynamically expanded when the QBUFHITH parameter value is reached. The default percentage is 20%.

**QBUFSZ=**

The size of the queue buffers.

**SHAREDQ=**

The suffix of the shared-queues PROCLIB member data set. This parameter is in the PARMLIB member DFSSQxxx.

**SHMSGSZ=**

The record length of the short message.

Define the following IMS startup parameters in the PROCLIB member data set DFSSQxxx:

**CQS=**

The PROCLIB member data set that contains the CQS procedure. CQS= is optional. The default is CQS=CQS.

**CQSSSN=**

The name of the CQS address space subsystem.

**EMHQ=**

The name of the EMH queues structure. If this statement exists, an EMHQ structure (and the other associated structures) is required in the shared-queues environment.

**MSGQ=**

The name of the message queues structure.

**SQGROUP=**

A one- to five-character identifier that represents the z/OS cross-system coupling facility IMS shared-queues group name, which is concatenated to the letters DFS. All IMS systems that share the same set of structures must specify the same SQGROUP= name. The SQGROUP= name can be the same as the CQSGROUP= name, which is specified in the CQSIPxxx PROCLIB member data set.

**WAITRBLD=Y | N**

Specifies whether activity on the EMHQ structure should wait until CQS completes the structure rebuild process. WAITRBLD=NO specifies that activity on the EMHQ structure should continue while CQS rebuilds the structure. The WAITRBLD parameter is optional. The default is WAITRBLD=N.

**Related reading:** For more information on these IMS parameters, see *IMS Version 11 System Definition*.

---

## Using the Common Queue Server

IMS uses the CQS interface to manage the shared queues in a sysplex environment. From each IMS that acts as a CQS client, IMS uses this CQS interface to access the shared queues.

**Restriction:** When using the same LTERM name on multiple systems in a shared queues sysplex, all conversational and response mode reply messages must be processed by the originating terminal's system. This means that while a conversation or response mode operation is still in-progress on one IMS, the same LTERM name cannot be used, simultaneously or serially, by a terminal on another

IMS. This Transaction Manager (TM) restriction applies to both Fast Path and non-Fast Path in a shared-queues environment.

## Starting CQS

You can activate CQS in one of two ways:

- As a z/OS task, using the z/OS START command
- As a z/OS batch job

In addition, IMS automatically starts CQS, when appropriate.

## Restarting CQS

Depending on whether a CQS structure contains data, you can warm start CQS or cold start CQS:

- If the structure is empty, you must cold start CQS.
- If the structure contains data, you can either warm start or cold start CQS.

When you have completed restarting CQS, the CQS ready message is issued (CQS0020I).

### *CQS warm start*

During a warm start, CQS reads the checkpoint data set to find the log token representing the last system checkpoint. When CQS finds this log token, it initiates a warm start. If CQS fails to find this log token in the checkpoint data set, it reads the log token from the structure. If CQS finds the log token, it issues a WTOR in order to allow you to confirm the use of this token.

### *CQS cold start*

When CQS cold starts after connecting to an empty structure, CQS purges all log records for the structure and performs a system checkpoint. After the system checkpoint is complete, the structure and CQS restart is complete.

## Using CQS user-supplied exit routines

You can use the following exit routines to monitor and modify CQS activities:

- CQS Initialization/Termination exit routine
- CQS Client Connection exit routine
- CQS Queue Overflow exit routine
- CQS Structure Statistics exit routine
- CQS Structure Event exit routine

## Changing the size of coupling facility structures

The initial size of a structure on the coupling facility is determined by the value of the INITSIZE parameter in the CFRM policy. CQS allows you to dynamically reconfigure the size of a structure.

When the first CQS connects to a structure, the value specified for INITSIZE is the size of that structure. If enough free space does not exist for this INITSIZE value, the size of the structure becomes that of the available space in the coupling facility.

## Initiating system checkpoint

As a system checkpoint for recovering CQS information during restart, each CQS writes its own control blocks to the z/OS log. To retrieve this restart information, CQS also writes information to the CQS checkpoint data set. CQS does not quiesce activity while the checkpoint is in progress.

### *Checkpoint data sets*

For each structure pair, CQS maintains a checkpoint data set. CQS writes to a checkpoint data set and then uses it during restart. Checkpoint data sets are dynamically allocated during CQS initialization.

### *How CQS restarts after system checkpoint*

During CQS restart, CQS reads the log records from the last system checkpoint and rebuilds the work that was in progress.

## Initiating structure checkpoint

In order to recover message queues, CQS takes structure checkpoints. CQS copies the shared queues from a structure pair to a structure recovery data set (SRDS). CQS quiesces activity while it performs part of this copy operation. CQS then performs a system checkpoint following each structure checkpoint.

## Recovering structures

z/OS allows you to rebuild structures. You rebuild structures either by copying them or by recovering them.

### *Rebuilding structures*

One or more CQs must be running in order to copy or recover structures and the messages on those structures. When the new structure is allocated, policy changes (such as structure location) are applied.

### *Copying structures*

If at least one CQS has access to the structure when structure rebuild is initiated, one of the CQS systems that still has access to the structure copies all of the messages (both recoverable and irrecoverable) from that structure to the new structure.

### *Recovering structures*

If no CQS has access to the structure when structure rebuild is initiated, the structure is recovered from the SRDS and the z/OS log. Irrecoverable messages (such as Fast Path input messages) are lost.

## Deleting structures

You can delete a structure when no CQS is connected to it. To delete a structure:

1. Shut down all CQs that are connected to the structure.
2. Delete any failed persistent connections.

**Attention:** If a CQS fails while it is connected to a structure, allow it to restart so that it can clean up any work that was in process at the time it failed. This command can be used to terminate the failed connections when you must delete the structure. If this command is used incorrectly, the queues or resources might be lost.

Issue the following command:

```
SETXCF FORCE,CONNECTION,STRNAME=strname,CONNAME=ALL
```

3. Issue the command:

```
SETXCF FORCE,STRUCTURE,STRNAME=strname
```

Ensure that the *strname* value in this command is the same as the *strname* value that is specified in the CQS global structure definition parameters in PROCLIB member data set CQSSGxxx and the CQS local structure definition parameters in PROCLIB member data set CQSSLxxx.

## Deleting message queues

You can delete message queues on a structure by deleting the structure.

## CQS logging and the z/OS logger

The z/OS system logger records all information necessary for CQS to recover structures and restart. CQS provides a File Select and Formatting Print utility to read the log records.

**Related reading:** For more information about defining the log stream, see *z/OS MVS Programming: Sysplex Services Guide*.

**Related reference:**

➡ CQSSGxxx member of the IMS PROCLIB data set (System Definition)

➡ CQSSLxxx member of the IMS PROCLIB data set (System Definition)

➡ File Select and Formatting Print utility (DFSERA10) (System Utilities)

## Monitoring and requeuing structures using the Queue Control Facility

You can copy and remove messages, and query shared queues for message counts and ages, using Queue Control Facility (QCF) two ways: by submitting a QCF BMP job to run on an IMS in the shared-queues group, or by using the QCF TCO/ISPF interface.

Removing messages can be useful for clean-up or for re-inserting and processing messages later. You can also select and re-queue lost messages to the shared queues for reprocessing.

## Requeuing the cold queue using QCF

When an IMS is abended and then cold started, CQS places messages in progress at the abend on the cold queues of the CQS shared-queues structure. You can analyze or delete these messages on the cold queue structure using QCF. Additionally, with QCF, you can move the messages back to one of the processing queues for re-processing.

**Related reading:** For more information on the QCF tool, see *IMS Queue Control Facility for z/OS User's Guide*.

---

## Accessing CQS with IMS commands

These IMS commands apply to the local IMS only. Commands and command responses are not placed on the shared queues; therefore, they are not recoverable.

You can use the following IMS commands to interact with a CQS:

**/CQCHKPT SHAREDQ**

Initiates a CQS structure checkpoint.

**/CQCHKPT SYSTEM**

Initiates a CQS system checkpoint.

**/CQQUERY STATISTICS**

Displays statistics for primary and overflow structures.

**/CQSET SHUTDOWN SHAREDQ**

Sets status to take structure checkpoint at CQS shutdown.

**/DEQUEUE**

Dequeues one or more messages from one of the following:

- An outbound APPC queue
- A Fast Path program
- An LTERM
- An MSNAME
- An outbound OTMA queue
- A message queue
- A remote transaction or a remote LTERM
- A transaction

**/DEQUEUE SUSPEND**

Dequeues one or more suspended transactions.

**/DISPLAY CQS**

Displays CQS information that IMS tracks.

**/DISPLAY MODIFY**

Displays local work in progress.

**/DISPLAY OVERFLOWQ**

Displays a list of queue names that are in overflow mode.

**/DISPLAY QCNT**

Displays global queue information for a particular resource type.

**/DISPLAY STRUCTURE**

Displays structure status.

**/DISPLAY TRACE TABLE QMGR**

Displays queue manager trace table status.

**/DISPLAY TRACE TABLE SQTT**

Displays shared-queues trace table status.

**/TRACE SET OFF TABLE QMGR**

Stops trace of queue manager activity.

**/TRACE SET ON TABLE QMGR**

Starts trace of queue manager activity.

**/TRACE SET OFF TABLE SQT**

Stops trace of shared-queues activity.

**/TRACE SET ON TABLE SQT**

Starts trace of shared-queues activity.

**QUERY TRAN**

Displays information about messages on the shared queues.

**QUEUE TRAN**

Enqueues and dequeues messages on the shared queues.

---

## List structures

CQS manages message queues that are shared within an IMSplex by using *coupling facility list structures*. A *primary list structure* holds the message queues. An *overflow list structure*, if you define one, holds additional queues after the primary list structure reaches a predefined threshold.

For IMS systems that do not share Fast Path transactions, define only message queue primary structures and optional message queue overflow structures. For IMS systems that do share Fast Path transactions, define both message queue primary (and optional overflow) structures and expedited message handler queue (EMHQ) primary (and optional overflow) structures.

Figure 29 on page 225 and Figure 30 on page 226 show shared queues for IMS clients A and B.

In the following figure, only message queue structures, message queue structure recovery data sets (SRDSs), message queue checkpoint data sets, and a message queue z/OS log stream are used.




Figure 29. Shared-queues environment with its message queue (MSGQ) list structures

The configuration shown in the following figure shows EMHQ structures, along with their corresponding SRDSs, checkpoint data sets, and a z/OS log stream. A configuration such as this one would also include message queues, but these are not shown in the figure.






Figure 30. Shared-queues environment with its EMHQ list structures

## Using associated printers

In a shared-queues environment, when a user signs on and specifies associated printer information, the IMS front end registers interest for each associated LTERM. When the application inserts a message to one of these LTERMs, the IMS front end is notified and begins autologon for the printer.

**Recommendation:** Do not provide autologon parameters for these users in the IMS back end; doing so causes both the front end and back end to process autologon for the associated printers.

For shared printers in a shared-queues environment, shared printer terminals are often switched between IMS systems in the IMSplex. This enables the output to be sent to any IMS in the IMSplex and to be printed.

---

## Planning for IMS front-end switch

Using IMS Front-end Switch (FES) in a shared-queues environment requires one of two events.

To complete an IMS front-end switch, either:

- Coding the FES exit routine (DFSFEBJ0) to route the input messages to the correct LTERM (to deliver FES reply messages correctly).
- Establishing ISC sessions between IMS systems within the IMSplex (to ensure that FES input and reply messages are sent and received by any IMS).

### **Related reading:**

- For information on the FES exit routine, see *IMS Version 11 Exit Routines*.
- For information on establishing ISC sessions, see *IMS Version 11 Communications and Connections*.

---

## Determining eligibility for sysplex-wide processing

Some IMS messages in a shared-queues environment are eligible for sysplex-wide processing; others are not.

**Definition:** An *IMS message* can be one of the following statements:

- A transaction or transaction response
- A message switch
- A system message
- A message from LU 6.2 or OTMA
- A command or command response

All messages (not including commands or command responses) are placed in shared queues.

### **Messages eligible for sysplex-wide processing**

Messages for the following types of communication are eligible to be processed on any IMS in the IMSplex:

- Transactions defined locally or using the Destination Creation exit routine (DFSINSX0)
- Fast Path transactions defined using the DFSHAGU0 exit routine with sysplex processing option LOCAL FIRST or GLOBAL ONLY
- Fast Path messages for a program that is active in the IMSplex
- The following synchronous message types:
  - APPC
  - OTMA
- LTERMs or by using the Destination Creation exit routine (DFSINSX0)

**Definition:** A message for an LTERM can be:

- A transaction response
- A message switch
- A system message
- A command response

## Messages not eligible for sysplex-wide processing

Messages for the following types of communication must be processed locally and are therefore not eligible for sysplex-wide processing:

- Serial transactions (transactions defined with SERIAL=YES specified on the TRANSACT macro)
- Commands and command responses

---

## Managing APPC and OTMA messages in a sysplex environment

APPC and OTMA messages can be either asynchronous or synchronous.

### Asynchronous APPC and OTMA messages

To ensure delivery of APPC and OTMA messages in a sysplex environment, enable APPC and OTMA on every back-end IMS in the shared-queues group. If a back-end IMS does not have APPC or OTMA enabled, any asynchronous APPC or OTMA output that is inserted to an alternate PCB is queued and not delivered until the operator issues a /STA APPC or /STA OTMA command.

Asynchronous APPC and OTMA messages created by a program-to-program switch can run on any IMS in the shared-queue environment.

### Synchronous APPC and OTMA messages

Synchronous APPC and OTMA (send-then-commit) messages can run on any IMS in the shared-queues environment to distribute the transaction workload. Synchronous APPC or OTMA output inserted into the I/O PCB must be delivered by the front-end IMS. Therefore, APPC or OTMA route the synchronous output back to the front-end IMS, regardless of which IMS is running the transaction. Non-conversational I/O PCB reply messages (less than 61 KB) are sent to the front-end IMS using z/OS cross-system coupling facility services. Conversational I/O PCB reply messages or any synchronous output messages greater than 61 KB are sent to the front-end IMS using shared queues and a special NOTIFY message that is sent using XCF. The IMSplex supports synchronous APPC and OTMA messages when the following conditions are true:

- All IMS systems in the shared-queues group are Version 9 or later.
- If AOS=Y, z/OS Resource Recovery Services is active.
- If AOS=Y, the IMS control region parameter RRS=Y is defined for all the IMS systems in the IMSplex.
- AOS=Y or AOS=F is specified in the DFSDCxxx PROCLIB member data set.

If a synchronous APPC or OTMA transaction running in a back-end IMS results in asynchronous APPC or OTMA output that is inserted to an alternate PCB, APPC or OTMA must be enabled on the back-end IMS. Also, the asynchronous output is delivered to the APPC or OTMA client directly from the back-end IMS.

Synchronous APPC and OTMA transactions initiated by a non-conversational program-to-program switch to local transactions always run on the same IMS system as the transaction that initiated them. For example, if synchronous TRAN A, which can run on any front-end or back-end IMS system in the sysplex, creates synchronous TRAN B and asynchronous TRAN C, both TRAN B and TRAN C will run on the same IMS system as TRAN A.

| In a shared-queues environment with APPC/OTMA SMQ enablement active, a  
| transaction coming into IMS through APPC/OTMA can be executed by any  
| member in the IMSplex. But when the transaction issues a program-to-program  
| switch, the switch to transaction must be processed by the same IMS member  
| where the switch occurred unless the control region parameters APPCASY=S or  
| OTMAASY=S is specified.

| In the APPC/OTMA shared queues environment when a transaction is processed  
| at the back-end IMS system and performs multiple program to program switches,  
| whether a switch to transaction can be scheduled synchronously is determined at  
| the GU IOPCB time. However, if the control region parameter APPCASY=S or  
| OTMAASY=S is specified, the decision as to which transaction is scheduled  
| synchronously get made at ISRT ALTPCB time.

**Related concepts:**

“Specifying APPC generic resource names” on page 291

 Administering APPC/IMS and LU 6.2 devices (Communications and Connections)

---

## Managing the IMS dead-letter queue

IMS maintains a dead-letter queue only for local user structures. For shared queues on the coupling facility, use the /DIS QCNT MSGAGE command to display potential dead-letter queues.

By examining the display output, you can determine whether a particular queue is a dead-letter queue and then take appropriate action.

---

## Using the Expedited Message Handler

The IMS Fast Path Expedited Message Handler (EMH) can also use the shared queues in a shared-queues environment. IMS calls the Fast Path Input Edit/Routing exit routine to determine if an incoming message should be processed by Fast Path.

If Fast Path must process the message, either the local IMS system processes it, or IMS places it on the shared queues.

Fast Path can use EMHQ structures, which are defined in the same way as the MSGQ structures (in the DFSSQxxx member of PROCLIB), to share Fast Path transactions. However, an IMS system with Fast Path installed is not required to have an EMHQ structure or share Fast Path transactions. If you do not define an EMHQ structure, all Fast Path transactions are processed locally.

If an IMS system joins the z/OS cross-system coupling facility group, abends, or shuts down, it notifies all sharing IMS systems by sending a message with the program name table. Sending this message makes all IMS systems aware of which IMS systems are servicing particular programs. Likewise, when a program is started or stopped on an IMS system, it notifies all sharing IMS systems, and each IMS system updates its program name table.

If no active region is available in the IMSplex to service a program, all IMS systems that have enqueued messages for that program issue message DFS2529I to the inputting terminals. These terminals are then unlocked. Any new input for the program is rejected with message DFS2533I.

Because Fast Path input messages are not recoverable, they are not written to the SRDS data set during CQS checkpoint. You can recover EMHQ structures; however, no Fast Path input messages are retained on the SRDS data set, and only Fast Path output messages can be recovered from the EMH queues. For each terminal awaiting output from an input message that has been lost during EMHQ structure rebuild, IMS sends message DFS2766I and unlocks the terminal. In addition, IMS writes a X'5936' log record for the lost message.

The EMHQ structures can be copied, and all messages are copied from the old structure to the new one, regardless of whether they are recoverable.

For EMHQ rebuild, it is possible to specify that the rebuilt EMHQ structure be accessible while the rebuild is in progress. Use the WAITRBLD= parameter in the DFSSQxxx PROCLIB member data set.

---

## Requeuing suspended transactions

When a transaction is suspended, it is placed on the transaction-suspend queue in the coupling facility. To make such messages available for processing, issue the /DEQUEUE SUSPEND TRAN command. Issuing this command moves the transaction to the transaction-ready queue or the transaction-serial queue.

The IMS that places the transaction on the transaction-suspend queue can be different from the IMS that issues the /DEQUEUE SUSPEND TRAN command. In this case, other IMS systems that share the message (MSGQ) structure need to issue the /DEQUEUE SUSPEND TRAN command.

**Related reading:** For more information on the transaction-suspend queue, the transaction-ready queue, or the transaction-serial queue, see “Queue types” on page 206.

---

## Scheduling AOI transactions

To ensure that the transaction runs locally, define all AOI transactions that are scheduled by the AOI exit routine as serial. Placing the AOI transaction on the shared queue allows any IMS to process it, regardless of whether the command applies to that IMS or not.

---

## Planning for MSC in a shared-queues environment

IMS systems in a multiple systems coupling (MSC) network can be part of a shared-queues environment, connected to IMS systems outside the shared-queues environment.

Each IMS within the IMSplex can act as any of the following:

- A front end, receiving messages from terminals or across MSC links
- A back end, processing messages received from a front end
- An intermediate subsystem, receiving messages and passing them on to other IMS systems across MSC links

Although two IMS systems can have MSC links defined, messages are placed on the shared queues from one of these IMS systems and picked up for processing by another IMS within the IMSplex—these messages are not sent over the MSC links.

**Exception:** If one of these IMS systems is running without shared queues enabled, the MSC links are used. This situation might occur when migrating to a shared-queues environment.

To plan for using MSC in a shared-queues environment, take each of the following actions:

- Specify all remote IMS systems on the MSNAME macro (SYSID and NAME), and assign a unique SYSID within the local IMS. These remote MSNAMEs remain stopped and are used only to route messages from any IMS within the IMSplex to a remote IMS.
- Define MSC remote system identifiers (SYSIDs), using the MSNAME macro, for each IMS within the IMSplex.
- Within the local IMS, make all SYSIDs be unique across all IMS systems in the IMSplex and in the MSC network.
- Define remote transactions and LTERMs to the IMS that has the MSC link (the back-end IMS). Also define them to the front-end IMS. If you do not define them, IMS calls the Destination Creation exit routine (DFSINSX0) to determine whether the destination is local or remote, and to determine whether it is a transaction or an LTERM.

When an MSC link exists on the front-end IMS, you can use MSC for APPC and OTMA transactions in a shared-queues environment. The conversation is maintained with the local IMS but is not carried to the remote IMS. Therefore, the remote application program cannot issue the SETO DL/I call, or the CPI-C verbs SEND\_ERROR and DEALLOCATE\_ABEND. Also, if the remote IMS allocates another APPC conversation, it receives the default user ID that is associated with the MPP region, and it does not use the original APPC user ID in order to verify security.

---

## Planning for RSR in a shared-queues environment

Remote Site Recovery (RSR) allows you to quickly recover computer services in the event of an interruption. In a non-shared-queues environment, RSR recovers full-function databases, Fast Path DEDBs, IMS message queues, and the TM network.

RSR is supported in a shared-queues environment.

**Recommendation:** All IMS systems that share a structure should:

- Be part of the same RSR service group
- Send log data to the same IMS tracking subsystem

If an RSR takeover occurs, and the active subsystems are restarted at the remote site, the new active subsystems must be started with the shared-queues option.

**Restriction:** For RSR in a shared-queues environment:

- You cannot recover the shared message queues or the TM network after a remote takeover, because message queues are not logged in the IMS log. After takeover, cold start the DC system using the /ERESTART command.
- The tracking subsystem cannot track messages in real time, because it cannot run with the shared-queues option. To handle local message activity, the queue manager message queue data sets are used.

---

## Tuning for performance in a shared-queues environment

The structures and parameters you define for a shared-queues environment determine how the system performs.

To optimize system performance, try adjusting the following structures and parameters:

- The size of the IMS execution parameters:
  - LGMSGSZ=
  - QBUF=
  - QBUFHITH=
  - QBUFLWTH=
  - QBUFMAX=
  - QBUFPCTX=
  - QBUFSZ=
  - SHMSGSZ=
- The value of the CQS local structure definition parameter SYSCHKPT=
- The values of the CQS global structure definition parameters:
  - OVFLWMAX=
  - STRMIN=
- The size of the shared-queues structures on the coupling facility
- The size of the z/OS system log structure on the coupling facility
- The number of z/OS system log data sets that back up the z/OS system log structure
- The frequency of structure checkpoints

### Recommendation for Fast Path transactions in a shared-queues environment

For shared Fast Path transactions, if you specify LOCALONLY in the Fast Path Input Edit/Routing exit routine (DBFHAGU0), do not change the size of your EMH buffer pools. The same buffer pool space is required for a shared-queues environment as for a non-shared-queues environment.

However, if you specify LOCALFIRST in DBFHAGU0 for shared Fast Path transactions and the transaction is processed locally, both a front-end EMH buffer and a back-end EMH buffer are required in the local IMS system; consequently the EMH buffer pool usage increases.

In either case, analyze EMH buffer usage and adjust the space as needed.

**Related reading:** For more information on tuning for performance in a shared-queues environment, see “Planning for performance in a shared-queues environment” on page 426.



---

## Chapter 15. Data sharing in IMS environments

When more than one IMS system has concurrent access to an IMS database, those systems are sharing data, and you have a data-sharing environment. In a data-sharing environment, you manage multiple IMS systems, each running on different z/OS systems in a sysplex. Such configurations require more sophisticated operations and recovery procedures.

This information describes how to allow more than one IMS online or batch system concurrent access to data that resides in a common database. If you use IMS to control this common access to data, the data sharing support provided by Database Recovery Control (DBRC) is required.

---

### Data sharing overview

An IMS system includes databases whose data can potentially be made available to, or shared with, all declared application programs. Access to a database is a characteristic defined in an application program's PSB. With data sharing support, application programs in separate IMS systems can concurrently access databases.

IMS systems use lock management to ensure that database changes at the segment level originating from one application program are fully committed before other application programs access that segment's data.

Data sharing among IMS systems is supported in both sysplex and nonsysplex environments.

- *Sysplex data sharing* is data sharing between IMS systems on different z/OS operating systems. IRLM uses a coupling facility to control access to databases.
- *Nonsysplex data sharing*, also referred to as *local data sharing*, is data sharing between IMS systems on a single z/OS operating system. A coupling facility can be used, but it is not required.

With data sharing, two levels of locking control are possible:

- *Database-level sharing*, in which an entire database is locked while an application program is making updates. Locking prevents concurrent database access and the scheduling of application programs that might jeopardize database integrity. For DEDB area resources, this is called *area-level sharing*.
- *Block-level sharing*, in which you can use global block locking to maintain database integrity during concurrent access of a database. The blocks are locked rather than the entire database. Multiple application programs can update a database at the same time if they are updating different blocks.

Within a single z/OS operating system, or in a sysplex, multiple IMS online or batch systems can update a database concurrently. Within a database, resources are reserved at the block level. For OSAM databases, the block is a physical block stored on a direct access storage device. For VSAM databases and DEDBs, the block is a control interval (CI). Data integrity is preserved for the multiple IMS online or batch systems that concurrently access the shared data. Sequential dependent segments are supported. *Block-level sharing* for DEDBs is between IMS online systems only.

Data access is protected so that:



- One IMS system is authorized to update the database, and other authorized IMS systems can have read-only access.
- Multiple authorized IMS systems can concurrently schedule the database with read or read-only access.

In a data-sharing environment, an IMS system can be online or batch. For area-level sharing, participating IMS systems must be online. IMS utilities are considered as batch systems that work with database-level sharing.

Some differences exist in support for data sharing configurations. Generally, a complete database is regarded as one data resource. When the database is invoked within an IMS online system, or as a batch IMS system, it must be available for an individual application program to process. However, a database might not be available if, for example:

- The database is used exclusively by one IMS system.
- The database is flagged as needing recovery.
- Backup procedures are in process.

With data entry databases (DEDBs), however, the data resource is divided. In a DEDB, each individual area is considered a unit of data resource. In this topic on data sharing, when we use the term “database,” it means a DEDB area, unless otherwise noted.

With High Availability Large Databases (HALDBs), the database can be divided into one or more partitions. In a HALDB database, each individual partition is considered a unit of data resource. When the term “database” is used in reference to data sharing, it means a HALDB partition, unless otherwise noted.

Restrictions associated with data sharing are:

- Batch IMS support excludes the use of MSDBs and DEDBs.
- Only IMS online systems that use Fast Path can share DEDBs.
- Data sharing support excludes MSDBs and GSAM databases.

## DBRC and data sharing support

Concurrent access to databases by systems in one or more z/OS operating systems is controlled with a shared Database Recovery Control (DBRC) RECON data set. IMS systems automatically sign on to DBRC, ensuring that DBRC knows which IMS systems and utilities are currently participating in shared access.

Subsequently, a system's eligibility to be authorized to access a database depends on the declared degree of sharing permitted and other status indicators in the RECON data set.

To maintain data integrity, status indicators in the RECON data set control concurrent access and recovery actions for the databases. This common RECON data set is required in a data sharing IMSplex because a given database must have a DMB number that uniquely identifies it to all the sharing IMS systems. The DMB number that DBRC records in its RECON data set is related to the order in which databases are registered to DBRC. Using multiple RECON data sets can result in the same DMB number existing in each RECON data set for different databases. This condition can result in damage to databases.

## Registering with DBRC

Databases that are to take part in data sharing must be registered in RECON. Each registered database has a current status that reflects whether it can take part in sharing and the *scope* of the sharing. The concept of scope combines several ideas:

- The type of access: read or update
- Whether more than one access can occur within the database simultaneously
- Whether an IMS system that needs access is in the same or a different z/OS operating system

You specify the level of sharing on an individual IMS database using the SHARELVL specification on the INIT.DB command.

Databases can be shared at the database level, at the block level, or not at all. The type of sharing used for a database should be based on the data integrity and availability needs of all the application programs using it.

When two or more IMS systems are executing concurrently and sharing data at the database level, they can be in the same or different z/OS operating systems.

Figure 31 on page 236 shows a database as a shared resource and also illustrates area-level sharing. One or more areas making up the DEDB could be a shared resource. The RECON data set is also shared as a VSAM key-sequenced data set (KSDS), with status being read and updated by a DBRC control portion of the online or batch IMS. For IMS online systems, DBRC executes in a separate region, automatically started by the control region at initialization. For a batch IMS, DBRC control is contained within the batch region, as it is for database-level sharing.

The RECON data set keeps track of the:

- Sharing level allowed for each database
- Databases or areas currently authorized for processing
- IMS systems that are involved
- Status of all of those systems
- Database status from a recovery viewpoint




Figure 31. Database-level sharing

## Naming conventions for data-sharing environments

If you decide to use data sharing and have additional databases to control with DBRC, you must review the data set naming conventions established for database data sets.

Consider indicating that a database participates in data sharing using the chosen resource names. An application program that accesses a shared database might also have the PSB name indicate a read-only property.

Review naming conventions for the following resources:

- Databases, their DD names and data set names
- Image copy and change accumulation data set names
- Online log data sets (OLDSs)
- System log data sets (SLDSs)
- PSB and application program names
- Transaction codes
- Identifying names for IMS systems and IRLMs

## Database integrity without data sharing support

Without data sharing support, the responsibility for database integrity lies with administration and operations.

Two situations where data integrity is not protected are:

- In a single z/OS operating system, with the database JCL specifying a disposition of SHR, multiple IMS online or batch systems executing simultaneously might concurrently access the database without integrity.
- If multiple IMS systems are executing in more than one z/OS operating system, and a database is on shared DASD, those IMS systems can concurrently access the database. This concurrent access cannot be directly controlled by JCL, even when a disposition of OLD is specified.

---

## How applications access data

Application programs can access data in three ways: update, read, or read-only access.

You indicate the type of access on the PROCOPT keyword as part of the group of statements that comprise the program communication block (PCB) for a particular database access. These processing options for an application program are declared in the program specification block (PSB) and express the data access and alteration intent of the application program.

If an application program is to insert, delete, replace, or perform a combination of these actions, the application program has *update access*, which is specified with the PROCOPT=A keyword. An online application program that has exclusive access is also interpreted as having update access; update access is specified with the PROCOPT=E keyword.

Application programs that need access to a database but do not update the data can do so in two ways. They can access the data with the assurance that any pending changes have been committed by the application program that initiated the change. This type of access is *read access* and is specified with the PROCOPT=G keyword. Alternatively, application programs can read uncommitted data, if the application program does not specify protection of data status. This type of access is *read-only access* and is specified with the PROCOPT=GO keyword.

For more information on PROCOPT keyword values, see *IMS Version 11 Application Programming*.

---

## How IMS systems share databases

This topic describes how database access is established, data sharing at both the block and database level, and how data access is controlled by each level of sharing.

### Establishing database access

To specify how an IMS system that is requesting access to a database plans to use the database, use the ACCESS keyword parameter in the DATABASE macro.

You use the ACCESS keyword in the /START command to show how the IMS system requesting access to a database plans to use the database. The presence of any update intent means the whole system requires update access for that database. If any database PCB requires exclusive use, the whole system requires update access for that database. If no updating application programs need update access, you must find out what kind of read access is required. If the online system is to use block-level sharing, you usually declare read access.

The access mode can be *update* (UP), *exclusive* (EX), *read* (RD), or *read-only* (RO).

If the IMS system requires exclusive use of a database, it can declare its access as *exclusive*. With exclusive access, the IMS system can insert, delete, replace, or perform a combination of these actions. Exclusive access, therefore, precludes data sharing in block-level or database-level configurations.

If the IMS system needs to insert, delete, replace, or perform a combination of these actions and must participate in data sharing, it can declare its access as *update*.

IMS systems that plan to read from a database, but not update the database, can use either read or read-only access. With *read access*, the IMS system can access the data with the assurance that any pending changes have been committed by the application program that requested the change.

With *read-only access*, the IMS system is allowed to read uncommitted data.

How you specify database access for an IMS system depends on whether you are running a batch or online system.

If you plan to use the online version of the Database Image Copy utility, you do not need to specify ACCESS=UP. Although the utility requires a somewhat limited access while it is creating an image data set, the access and authorization are managed by DBRC.

### **Declaring database access for IMS batch systems**

For an IMS batch system, the database access directly corresponds to the highest PROCOPT value specified in the application program's PCBs or SENSEG statements. For example, if one PCB has PROCOPT set to G, and another PCB for the same database has PROCOPT set to I, the value of I is used; that is, the access for that database is update.

### **Declaring and changing database access for online systems**

For an online IMS, you specify access with the ACCESS keyword in the DATABASE macro during system definition. The access is declared for each individual database and reflects a desired access suitable for all the application programs that might be scheduled against that database.

**Related reading:** For a detailed description of the ACCESS keyword, see *IMS Version 11 System Definition*.

To dynamically change how an online IMS system accesses an individual database, use the /START command, which is described in *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

---

## **Using IRLM with database-level sharing**

When IMS uses the Internal Resource Lock Manager (IRLM) for locking services, IRLM provides all locking services.

If IRLM was not activated with the IRLMNM parameter in the IMSCTRL macro, and you want to activate IRLM, code IRLM=Y in the execution JCL. If the IRLMNM parameter is not specified in either the IMSCTRL macro or the execution JCL, and IRLM=Y is specified in the execution JCL, the IMSCTRL macro uses the default IRLM named ('IRLM').

When all of the IMS batch and online instances use one or more communicating IRLMs with database-level sharing, database extensions and buffer invalidations are notified to the read-only IMS instances. This is the same as full block-level data sharing, but without the locking activity.

When you use IRLM for database-level sharing, application programs that read without integrity read valid, but perhaps uncommitted, data more frequently. This does not provide any guarantees that read without integrity problems are solved (see *IMS Version 11 Application Programming*), but it does improve database-level sharing. Because of the order in which all changed blocks or control intervals (CIs) are written to DASD, including when they are written, an application program reading without integrity can still experience inconsistencies from one execution to another.

IMS allows use of the coupling facility for buffer coherency for both database-level sharing and block-level data sharing.

## Block-level sharing

Block-level sharing requires the Internal Resource Lock Manager (IRLM). If each z/OS operating system contains an IMS online system participating in block-level sharing, an IRLM component must exist in each z/OS operating system.

IRLM manages all requests for locks that control access to resources participating in data sharing.

Each IRLM responds to the requests that originate from its associated IMS systems. To do this, it maintains records of the participating IMS systems, as well as the status of locks that are held and waiting.

When an IMS online system uses IRLM, IRLM also provides the lock management that controls database resources that are concurrently accessed by application programs within the online system. In this way, the IRLM associated with an IMS online system provides almost all locking services for that system.

The IRLMs communicate with each other and manage the status of locks held for different database locks.

When using IRLM, you must define DBRC-registered SHARELVL=1 VSAM databases to VSAM with the appropriate SHAREOPTIONS required for block-level data sharing. This might require you to make a VSAM definitional change.

## Multiple IRLMs in sysplex data sharing

Multiple IRLMs can take part in the control of data sharing at the block level when IMS systems are executing in the same z/OS operating system.

This could be the case when IMS online systems are being tested for their use of data sharing before one of the systems is installed on a second z/OS operating system.

A special case of sysplex data sharing occurs when multiple IRLMs execute on a single z/OS operating system. In this case, you declare global sharing and give each IRLM a unique z/OS subsystem name and an identifying IRLM number (ID).

---

## Setting up IRLM procedures

For each IRLM component that can be activated, you need a procedure that is invoked when the system console operator enters the z/OS START command.

An example of one such procedure is DXRJPROC, described in *IMS Version 11 System Definition*. If you plan to use more than one IRLM procedure on a single operating system, include additional members in IMS.PROCLIB or SYS1.PROCLIB. Additional IRLMs require a unique identifier and any desired changes in parameter values. The procedure name cannot be the same as the IRLM subsystem names established for z/OS.

For each IRLM component, the procedure parameters specify:

- The component's identifying name
- The scope of its control
- Sysplex data-sharing information
- IRLM actions in failure situations
- The amount of virtual storage
- Performance factors

The PARM1= and PARM2= positional parameters for the region that executes in support of block-level sharing are shown in the following table.

Table 14. Categories and purpose of IRLM region parameters.

| Category         | Parameter (PARM1 and PARM2) | Purpose                                                                                                                                |
|------------------|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| Component name   | IRLMID                      | An identifier for this IRLM.                                                                                                           |
|                  | IRLMNM                      | z/OS subsystem name.                                                                                                                   |
| Scope of control | SCOPE                       | Inter-z/OS communication or intra-z/OS control.                                                                                        |
| Data sharing     | IRLMGRP                     | Name of the z/OS cross-system coupling facility group.                                                                                 |
|                  | MAXUSRS                     | Maximum number of users in the group.                                                                                                  |
|                  | LOCKTAB                     | Lock structure to be used by the group.                                                                                                |
| Storage          | MAXCSA                      | Maximum amount of z/OS common service area (CSA) to be used. This parameter does not apply to IRLM 2.2 and above.                      |
|                  | PC                          | Specifies control blocks in private storage. IRLM 2.2 always runs with PC=YES.                                                         |
|                  | PGPROT                      | Determines whether the IRLM load modules that are resident in common storage are placed in MVS PAGE PROTECTED STORAGE, default is YES. |
| Performance      | DEADLOK                     | Deadlock detection timing.                                                                                                             |
| Tracing          | TRACE                       | Trace types DBM, XCF, and SLM turned on at IRLM initialization (EXP, INT, and XIT are always on).                                      |

### Identifying IRLM

The IRLMID parameter specifies a decimal number (from 1 to 255) for IRLM identification. The IRLMID gets converted to a hexadecimal equivalent. A unique number must be assigned for each IRLM. IRLM error and status messages include the subsystem name and ID of the IRLM that issued the message.



The IRLMNM parameter specifies the 1- to 4-byte z/OS subsystem name that is to be assigned to this IRLM. Unless more than one IRLM is used at the same time on the same system, you can use "IRLM" as the subsystem name in each operating system. When IRLMs are used concurrently on a system, the subsystem names must be unique.

## Specifying IRLM scope

You must specify what kind of data-sharing control is required of the IRLM component—whether sysplex or nonsysplex sharing is to be used. The SCOPE parameter accomplishes this. SCOPE=GLOBAL specifies sysplex data sharing is to be controlled. This is sharing of resources among multiple systems.

SCOPE=LOCAL limits data sharing to nonsysplex data sharing. This is sharing of resources among IMS systems all identified to a single IRLM. If a coupling facility is not available, SCOPE=LOCAL is required. SCOPE=NODISCON specifies sysplex data sharing is to be controlled with special DISCONNECT rules. SCOPE=GLOBAL or NODISCON requires that a coupling facility be defined.

## Limiting use of CSA and ECSA by IRLM Version 2.1

The MAXCSA= keyword in the IRLM startup procedure specifies the maximum amount of common service area (CSA) and extended CSA (ECSA) storage space that IRLM Version 2.1 can use for its lock control blocks. You can specify the amount in bytes, kilobytes, or megabytes in the MAXCSA= keyword. For example, MAXCSA=5000000, MAXCSA=5000K, and MAXCSA=5M all specify that IRLM can use a combined total of 5 MB of CSA and ECSA storage space.

Use the PC= keyword in the IRLM startup procedure to reduce the CSA requirements. Specifying PC=YES causes the lock control blocks to reside in the IRLM 2.1 address space. In this case MAXCSA is ignored but must be non-zero. If IRLM 2.1 operates with PC=YES, the storage required for lock control blocks is placed in extended private storage. With PC=NO, the control blocks are in ECSA.

IRLM 2.2 always stores its lock control block structures in private storage and not in CSA and ECSA storage. Therefore, IRLM 2.2 always functions as if PC=YES were specified and ignores the PC= and MAXCSA= keyword values.

**Attention:** Both IRLM 2.1 and IRLM 2.2 use CSA and ECSA storage for other processing purposes not controlled by the PC= keyword.

## IRLM deadlock management

IRLM provides deadlock management. Using the DEADLOK parameter in the DXRJPROC procedure, you can specify the local deadlock-detection interval and the number of local cycles that are to occur before IRLM initiates a global deadlock detection. You can specify the local deadlock-detection interval in seconds or milliseconds. You can also dynamically modify the DEADLOK values using the z/OS command MODIFY IRLMPROC.

In a data-sharing environment, the participating IRLMs synchronize their DEADLOK parameters to the highest values provided. You should make this parameter the same for each IRLM in the data-sharing group.

### Related readings:



- For more information about the DXRJPROC procedure and the DEADLOCK parameter, see *IMS Version 11 System Definition*.
- For more information about the command MODIFY IRLMPROC, see *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*.

## System initialization for IRLM

To initialize the system for IRLM, you must define IRLM as a z/OS subsystem.

### Defining an IRLM as a z/OS subsystem

Your installation must assign a 4-byte subsystem name. You can use “IRLM” as the name, which assists the system operator in recognizing system messages sent to the console. This name must be unique to the z/OS system. If more than one IRLM subsystem is to run in an operating system, the names must be different.

An additional identification for each IRLM is a number in the range 1 to 255. The identification is required to internally differentiate between IRLMs. For example, two IRLMs that are to share data could have IDs of 1 and 2. Status and error messages sent to the system console operator contain the IRLM subsystem name and ID. Two IRLMs executing in a single z/OS system, or in different z/OS systems that are members of the same data-sharing group, require different ID numbers.

Another consideration for the IRLM subsystem name is that it must not be the same as that used for the startup procedure. This procedure is a member of SYS1.PROCLIB and is used by the operator in the START command. A unique START procedure must exist for each IRLM subsystem. The procedure library members are also distinguished by the different ID numbers and other control parameters.

The execution of the IRLM requires that the subsystem be added to the z/OS program properties table (PPT).

The system initialization routine for the subsystem must be null.

#### Related readings:

- For more information on defining the IRLM as a subsystem and adding an entry to the subsystem name table, see *z/OS MVS Initialization and Tuning Guide*.
- For more information on adding IRLM to the PPT, see “Adding the IRLM entry to the z/OS Program Properties Table” on page 129.

### Allowing for IRLM trace output printing

To provide a trace of the activity occurring in the IRLM subsystem, you use the z/OS Component Trace facility or the TRACE=YES option of DXRJPROC, the IRLM procedure.

### CTRACE records

IRLM produces trace output in z/OS component trace (CTRACE) format. This allows you to use IPCS CTRACE format, merge, and locate routines to process the buffer data. Both the IRLM trace buffers in the dump and external writer data set can be formatted using IPCS. The IRLM trace formatting load module (DXRRLFTB) and buffer find routine load module (DXRRL186) must be available for IPCS.

## Arranging for formatted dump output

IRLM uses the SDUMP facility of z/OS. The offline printing of dump output is done using z/OS print dump.

### Related reading::

- For more information on SDUMP, see *IMS Version 11 Diagnosis*.
- For information on DXRJPROC, the IRLM procedure, see *IMS Version 11 System Definition*.

## Defining an IRLM for sysplex data sharing

If you are implementing a sysplex data-sharing environment, you must define the data-sharing group to which IRLM belongs, the lock structure to be used by that group, and the maximum number of users in that group.

### Defining the data-sharing group

For an IRLM to belong to a data-sharing group, you must specify the name of the data-sharing group and the name of the lock structure in the IRLM startup procedure. All IRLMs in this group can share the same data. Each IRLM in the group must:

- Have a unique IRLMID
- Specify the same data-sharing group name using the GROUP parameter
- Specify the same lock structure using the LOCKTABL parameter

Although you can specify these definitions on the IRLM startup procedure, the recommended method is to define them using the CFNAMES control statement. The CFNAMES control statement defines the data-sharing group for a sysplex data-sharing environment. If you do not use the CFNAMES control statement, the lock structure name is pulled from the IRLM startup procedure.

### Defining the lock structure

Each IRLM participating in a sysplex data-sharing group (specifying the same z/OS cross-system coupling facility name on the GROUP=IRLMDS parameter) must specify the same lock structure. You specify the lock structure name using the LOCKTABL parameter.

### Defining the number of users in a data-sharing group

You must specify a maximum number of users for the data-sharing group. You do this by specifying a value from 2 to 32 on the MAXUSRS parameter.

---

## Data sharing at the database level with update activity

This type of sharing allows one IMS batch or online system to have update access. All other IMS systems must have read-only access.

For full-function databases, uncommitted data can be read. For Fast Path, uncommitted data does not exist, but all of the updates for a transaction might not yet have been written. This can cause invalid pointers to be referenced when Get Next processing is performed.

Normal application program isolation protects the integrity of the updating system. However, you can use IRLM for lock management on an IMS online system. The following figure illustrates an online system with update access sharing data with two batch systems that are authorized for read-only access to the same data.




Figure 32. Example of data sharing at the database level with update access

## Data sharing at the database level with multiple readers

This type of sharing allows multiple authorized IMS systems to read data concurrently with read or read-only access.

A sample configuration is shown in the following figure. Data integrity is not compromised, because no updates are allowed in this configuration.




Figure 33. Example of data sharing at the database level with read access

## Data sharing at the block level

Online IMS systems with full update capability can share data among themselves and can allow other batch IMS systems to read committed data or update the data.

The batch system can have the data integrity protected or use read-only access. This type of configuration is illustrated in the following figure.




Figure 34. Example of data sharing at the block level with update access

## Tailoring IMS systems that share data

Use the system definition macro and initialize the DBRC data sets and JCL to tailor an IMS system.

To tailor an IMS system to meet your data-sharing requirements:

- Use system definition macros to:
  - Include DBRC data-sharing support for IMS batch systems
  - Include IRLM (for block-level sharing)
  - Declare database access attributes
- Initializing DBRC data sets and JCL. See Chapter 41, “Initializing and maintaining the RECON data sets,” on page 535.
- Tailoring the execution JCL and virtual storage usage. See “Tailoring execution JCL” on page 248.

## Installation tasks

Before executing the IMS systems that use data sharing, you must complete several installation-related tasks.

The following table lists the tasks, indicates whether a task affects data sharing, and identifies other tasks that must be completed.

Table 15. Installation steps with additional data-sharing activity

| Step | Installation task                          | Affects systems without data sharing | Affects systems with data sharing | Additional task                                                                    |
|------|--------------------------------------------|--------------------------------------|-----------------------------------|------------------------------------------------------------------------------------|
| 1    | Build system libraries—IRLM                | No                                   | Yes                               | Add IRLM to system library.                                                        |
| 2    | Allocate and catalog IMS data sets         | Yes                                  | Yes                               | Prepare RECON data set.<br><br>Initialize DBRC controls.<br>Coordinate DL/I exits. |
| 3    | Prepare system definition and JCL          | Yes                                  | Yes                               | Macros and JCL changes.<br><br>Perform system definition.                          |
| 4    | Tailor the z/OS operating system           | No                                   | Yes                               | Add IRLM subsystem.                                                                |
| 5    | Tailor IMS performance options             | Yes                                  | Yes                               | Define buffers.                                                                    |
| 6    | Build DBDLIB                               | No                                   | No                                | None.                                                                              |
| 7    | Build PSBLIB                               | Yes                                  | Yes                               | Review PROCOPT values.                                                             |
| 8    | Build ACBLIB.                              | Yes                                  | Yes                               | Perform ACBGEN                                                                     |
| 9    | Selections for dynamic allocation          | Yes                                  | Yes                               | Review options.                                                                    |
| 10   | Prepare MFS libraries                      | No                                   | No                                | None.                                                                              |
| 11   | Prepare program libraries                  | Yes                                  | Yes                               | Review call sequences.                                                             |
| 12   | Perform initial database loading           | No                                   | No                                | None.                                                                              |
| 13   | Establish security                         | No                                   | No                                | None.                                                                              |
| 14   | Initialize IMS.MODSTAT                     | Yes                                  | Yes                               | Coordinate ACBLIB use and DBDLIB and PSBLIB.                                       |
| 15   | Copy staging libraries to active libraries | Yes                                  | Yes                               | Coordinate ACBLIB use.                                                             |

## DBRC and IRLM support for batch systems in a data-sharing environment

By default, DBRC provides data-sharing support for your batch and utility regions.

To change the default, you can either:

- Assemble and bind DFSIDEF0 into the IMS execution library (If you specify RMODE, you must specify RMODE=24 and AMODE=24)
- Override the value in a DFSPBxxx IMS PROCLIB member data set
- Override the value in JCL

To define the default IRLM support specifications for batch systems in a data-sharing environment, use the IRLM= keyword on the IMSCTRL macro. You can override specifications made in the IMSCTRL macro by using the IRLM= keyword in the batch JCL.

In batch systems, you can specify IRLM=Y or IRLM=N. If you specify IRLM=N:

- IRLM does not perform locking for the batch system.
- The batch system cannot participate in block-level sharing.

When batch systems do not use IRLM, DBRC still ensures the integrity of databases. DBRC authorizes batch jobs that have update access to a database only if all other IMS systems and batch jobs that are currently authorized by DBRC to the database have read-only access.

## Excluding a database from data sharing

If you specify ACCESS=EX, or leave the value for ACCESS= blank, the named database is not accessed concurrently by any other IMS system.

Register the database with DBRC, specifying a share level of zero.

Databases whose data does not need to be shared with any other IMS systems are owned by the IMS online system you define. If you do not want to use DBRC to control authorization to access these databases, do not register them.

---

## Tailoring execution JCL

The system log is required for any batch IMS that has update intent against any database and operates with active DBRC. The system log is not required if the batch system only uses read or read-only access intent.

When you are preparing a configuration that includes data sharing, you must:

- Specify system data sets.
- Specify the databases with DISP=SHR.
- Specify the placement of database data sets and device choices to suit the physical paths to the data.
- Prepare IRLM procedures.

The ACBLIB member online change process is coordinated among all IMS systems sharing the OLCSTAT data set. The system data sets that must be coordinated so that their status is common to all systems participating in data sharing are:

- RECON data sets
- Database data sets, with DISP=SHR specified
- IMS.ACBLIB, PSBLIB, and DBDLIB
- Libraries holding randomizing routines and DL/I exit routines

You also must set up IRLM procedures for your data-sharing configuration. For information on preparing the IRLM procedures for a data-sharing environment, see “Setting up IRLM procedures” on page 240.

---

## Tailoring the z/OS operating system

Because block-level sharing requires the use of IRLM installed as an independent component under the z/OS operating system, several tailoring actions must be performed.

You must:

- Coordinate VSAM data set definitions.
- Tailor the z/OS system for IRLM.
- Define an IRLM if you are sharing data across a sysplex.

## Coordinating VSAM data set definitions with share options

When your database access method is VSAM, the declaration of the data set indicates to VSAM what degree of shared access is required.

The SHAREOPTIONS parameter values (in the DFSMS DEFINE CLUSTER keyword) and the type of sharing they specify is shown in the following table.

Table 16. SHAREOPTIONS parameter specifications

| SHAREOPTIONS value | Type of sharing                                 |
|--------------------|-------------------------------------------------|
| (1,3)              | No sharing, single updater, or multiple readers |
| (2,3)              | Single updater and multiple readers             |
| (3,3)              | Multiple updaters and multiple readers          |

For databases that are to participate in block-level sharing and use the VSAM access method, you must include the SHAREOPTIONS (3,3) parameter when defining the data sets. The RECON data set is also accessed as a KSDS and requires SHAREOPTIONS (3,3).

**Related reading:** For preparation of DFSMS Access Method Services statements that declare and catalog the share options, see *z/OS DFSMS Access Method Services for Catalogs*.

When using IRLM, you must define SHARELEVEL=1 VSAM databases to VSAM with the appropriate SHAREOPTIONS required for block-level data sharing. This might require you to make a VSAM definition change.

### Tailoring the z/OS system for IRLM

Several other system initialization activities are required for the IRLM component:

- Defining the IRLM as a z/OS subsystem
- Allowing for IRLM trace output
- Arranging for formatted dump output
- Defining an IRLM for sysplex data sharing

These activities are described in “System initialization for IRLM” on page 242.

---

## Starting and stopping data sharing

Starting and stopping data-sharing operations involves starting and stopping the IRLM and database processing.

### Starting the IRLM

Start the IRLM from the system console of each participating system with the `START irlmproc,parms` z/OS command.

*irlmproc* is the unique name of the IRLM procedure, and *parms* represents one or more IRLM control parameters that override the command defaults.

Start the IRLM subsystems before starting the batch or online IMS.



If IRLM has not completed its initialization when the IMS system tries to connect to it, IMS issues a DFS039A message that says the IRLM is not active, and asks the operator to reply to the message: retry, cancel, or dump.

- If you have not started the IRLM, issue the z/OS START command, wait until the IRLM is active, and reply RETRY. IMS then tries to connect to the IRLM.
- If you have a problem coordinating startup, not merely starting the IRLM, reply CANCEL to cause the IMS system to terminate with abend U0039 .

**Example:** Assume the following environment:

- On processor J, an IMS system (IMSJ) uses an IRLM (JRLM).
- On processor K, an IMS system (IMSK) uses an IRLM (KRLM).

Next, assume the following scenario:

1. A hardware error occurs on processor K.
2. You try to restart IMSK on processor J to perform database backouts (but you do not change the IRLM name parameter, IRLMNM=KRLM).
3. Message DFS039A indicates that KRLM is not active.

Reply CANCEL in this case to immediately terminate the IMSK startup. Then, after changing IRLMNM= to JRLM, restart IMSK on processor J.

- If you reply DUMP, IMS terminates with abend U0039 and a storage dump.

The z/OS console operator must watch for messages originating from the IRLM. These messages have the three-letter prefix DXR.

## Stopping the IRLM

You can terminate an IRLM normally only when all IMS systems that are using it have terminated. To stop an IRLM normally, use the STOP *irlmproc* z/OS command.

To stop an IRLM abnormally, use the z/OS command MODIFY *irlmproc*,ABEND,NODUMP. To stop an IRLM abnormally and generate a dump, use the z/OS command MODIFY *irlmproc*,ABEND.

## Starting and stopping database processing

In a data-sharing environment, IMS type-1 commands can act *locally*, that is, affect the local system only, or can act *globally*, that is, affect all sharing IMS systems in the sysplex.

The UPDATE command can be routed to individual IMS systems or to all IMS systems in an IMSplex, using an automated operator program like the TSO SPOC.

An IMS command with the LOCAL keyword (or without the GLOBAL keyword) affects only the online IMS system on which you enter it. An IMS command with the GLOBAL keyword first affects the IMS system on which you enter it and then, if it is successful, affects other sharing IMS systems in the sysplex.

When you enter a /START DB or UPDATE DB START(ACCESS) command, either locally or globally, DBRC makes sure that each registered database does not need recovery or backout. The IRLM communicates global commands only in a block-level data-sharing environment.

The following table summarizes the commands that affect data sharing when issued for registered databases. A command entered for a database affects full function transactions; a command entered for an area affects Fast Path transactions.

Table 17. Database control commands across IMS systems

| Command                                         | Action by DBRC                                                                                                                                                                                  | Effect in local IMS system                                                                                          | Effect in sharing IMS systems           |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|-----------------------------------------|
| <b>Local</b>                                    |                                                                                                                                                                                                 |                                                                                                                     |                                         |
| /RMCHANGE                                       | Can change share level for IMS systems using the database                                                                                                                                       | None                                                                                                                | None                                    |
| /START AREA  <br>DATABASE<br>LOCAL or<br>UPDATE | Ensures that the database or area does not need recovery or backout, and records database start                                                                                                 | Allows transactions to access the database or area, and can change the access intent if you use the ACCESS= keyword | None                                    |
| <b>Global</b>                                   |                                                                                                                                                                                                 |                                                                                                                     |                                         |
| /DBDUMP<br>AREA  <br>DATABASE<br>GLOBAL         | Ensures that the database or area does not need recovery or backout, sets a return code for any unauthorized batch IMS systems, records database close, and changes access level to read access | Prevents online updating access                                                                                     | Processes command as if entered locally |
| /DBRECOVERY<br>AREA  <br>DATABASE<br>GLOBAL     | Prohibits further authorization                                                                                                                                                                 | Prevents allocation of database                                                                                     | Processes command as if entered locally |
| /STOP AREA  <br>DATABASE<br>GLOBAL              | Records database status as stopped (restricts further authorization)                                                                                                                            | Stops transactions from accessing the database or area                                                              | Processes command as if entered locally |
| /STOP ADS                                       | Records area data set as unavailable                                                                                                                                                            | Stops the area data set                                                                                             | Processes command as if entered locally |
| /START AREA  <br>DATABASE<br>GLOBAL             | Ensures that the database or area does not need recovery or backout, resets the prohibit-further-authorization and read-only flags, records database start, and can change the access level     | Allows transactions to access the database or area                                                                  | Processes command as if entered locally |

**Example:** You enter the /START command with the GLOBAL keyword on one IMS system and specify several database names, the IRLM transmits the command to other sharing IMS systems, deleting the names of any databases that are invalid for the local system before it transmits the command, and all sharing IMS systems process the command. In those online data-sharing systems, you see the DFS3334I message followed by the DFS3328I message, which tell you that the global /START command has started and has then completed. The messages include the database names that you include in your command.

If you omit the GLOBAL keyword (or specify the LOCAL keyword), the command applies only to the local online IMS system and does not affect access by any other IMS system.

When you enter a /STOP or UPDATE command locally, no interaction occurs between IMS and DBRC.

**Related reading:** For the format of these commands in a CICS environment, see *CICS Transaction Server for z/OS CICS Supplied Transactions*.

---

## Monitoring data-sharing systems

To monitor data sharing, you obtain information on the status of the following: IRLM, IMS systems and databases, the RECON data set, and coupling facility structures.

### Obtaining the status of IRLM activity

Use the MODIFY z/OS command to display the status of an IRLM.

To display the status of an IRLM on either your system or on another connected system, enter the z/OS command, where *irlmproc* is the name of the procedure that you used to start the IRLM, and *irlmx* is the name of the IRLM whose status you want to display:

```
MODIFY irlmproc,STATUS,irlmx
```

You can use the ALLD keyword to display the names and status of every IMS identified to an IRLM in a data-sharing group. Or, you can use the ALLI keyword to display the names and status of every IRLM in a data-sharing group.

You can also trace IRLM activity. see “Tracing IRLM activity” in *IMS Version 11 Operations and Automation*.

### Displaying components and resources

Monitoring components and resources in a data-sharing environment requires the same kinds of procedures as in a non-sharing environment.

The following table lists keywords for the /DISPLAY command that you can use to obtain information about various IMS resources.

*Table 18. /DISPLAY command keywords that provide information about IMS resources*

| Resources                                 | /DISPLAY command keywords                                                                                  |
|-------------------------------------------|------------------------------------------------------------------------------------------------------------|
| Active control regions                    | ACTIVE REGION                                                                                              |
| Active jobs                               | ACTIVE                                                                                                     |
| Programs, transactions, and conversations | CONVERSATION<br>PROGRAM<br>PSB<br>STATUS PROGRAM<br>STATUS TRANSACTION<br>SYSID TRANSACTION<br>TRANSACTION |
| Databases                                 | DATABASE<br>AREA<br>STATUS DATABASE                                                                        |

*Table 18. /DISPLAY command keywords that provide information about IMS resources (continued)*

| Resources                                                  | /DISPLAY command keywords                                                                                                                                                                                                      |
|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Terminals, lines, links, and nodes                         | ACTIVE DC<br>ASSIGNMENT LINE<br>ASSIGNMENT LINK<br>ASSIGNMENT NODE<br>LINE<br>LINK<br>LTERM<br>MASTER<br>MSNAME<br>NODE<br>PTERM<br>STATUS LINE<br>STATUS LINK<br>STATUS LTERM<br>STATUS MSNAME<br>STATUS NODE<br>STATUS PTERM |
| External subsystems and connections to external subsystems | CCTL<br>OASN SUBSYS<br>SUBSYS                                                                                                                                                                                                  |
| VTAM                                                       | TIMEOVER                                                                                                                                                                                                                       |

For example, you can issue the /DISPLAY DB command after you issue a /START DB command to determine whether the database is started.

The following table lists keywords for the QUERY command that you can use to obtain information about various IMS resources.

*Table 19. QUERY command keywords that provide information about IMS resources*

| Resources                                 | QUERY command keywords |
|-------------------------------------------|------------------------|
| Active control regions                    | MEMBER                 |
| Members in the IMSplex                    | IMSPLEX                |
| Transactions                              | TRAN                   |
| Databases                                 | DB<br>AREA             |
| Global online change status               | OLC                    |
| Language environment (LE) runtime options | LE                     |
| Online reorganization (OLR) status        | OLREORG                |
| RM resource structure                     | STRUCTURE              |

For example, you can issue the QUERY DB command after you issue an UPDATE DB START(ACCESS) command to determine whether the database is started.

**Related reading:** For detailed information about the /DISPLAY, /START, QUERY, and UPDATE commands, see *IMS Version 11 Commands, Volume 1: IMS Commands A-M* and *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

## Monitoring information in the RECON data set

Evaluate information from the RECON data set on a daily basis to understand the status of the total data-sharing workload and response times for all critical data-sharing transactions.

You can use the historical trace of events recorded in the RECON data set to find out the intervals when shared data was active. The LIST command or QUERY request can give you a comprehensive listing of events. For sysplex data sharing, you need a listing from all operating systems. Events such as IMS signon and time stamped database data set OPEN/CLOSE can be traced.

## Monitoring structures on a coupling facility

The DISPLAY XCF,STRUCTURE and DISPLAY XCF,STRUCTURE,STRNAME=z/OS operator commands are especially useful in monitoring structure activity.

These commands let you look at structures on a coupling facility to determine resource status and, for failures, gather information for problem determination.

For detailed information on these commands, see the appropriate z/OS publication.

### DISPLAY XCF,STRUCTURE command

Use this command to display the status of structures defined in your active policy.

The following example shows an example of output from this command.

```
IXC359I 11.09.26 DISPLAY XCF 376
STRNAME      ALLOCATION TIME  STATUS
CF01          02/17/96 17:03:49 ALLOCATED
CF02           --         --    NOT ALLOCATED
CF03           --         --    NOT ALLOCATED
CF04           --         --    NOT ALLOCATED
OSAMESXI      02/17/96 17:02:54 ALLOCATED
                                   REBUILDING
                                   REBUILD PHASE: QUIESCE
VSAMESXI      02/17/96 17:03:03 ALLOCATED
```

In the example:

#### STRNAME

Is the name of a structure.

#### ALLOCATION TIME

Is a time stamp indicating when the structure was allocated in the coupling facility.

#### STATUS

Is the current status of the structure. A structure can be in a number of different states, such as ALLOCATED, NOT ALLOCATED, or ALLOCATED REBUILDING.

The displayed information can help you determine whether the appropriate IRLM, OSAM, VSAM, shared MSGQ, and shared EMHQ structures have been defined and allocated. If not, check that the structure names on your CFNAMES control statement match the structure names that are displayed. If they do not match, change the names in either the coupling facility resource manager (CFRM) policy or the CFNAMES control statement.

In the example, six structures are defined (CF01-CF04, OSAMSESXI, and VSAMSESXI). Three of the structures are allocated (CF01, OSAMSESXI, and VSAMSESXI). The status of the OSAM structure is that it is currently being rebuilt after a structure failure.

## DISPLAY XCF,STRUCTURE,STRNAME= command

Use this command to display detailed information about a specific structure. The three structures of importance to IMS are the IRLM, OSAM, VSAM, shared MSGQ, and shared EMHQ structures.

The following example shows an example of output from this command, specifying an OSAM structure named OSAMSESXI.

```
IXC360I 17.30.46 DISPLAY XCF 677
STRNAME: OSAMSESXI
STATUS: ALLOCATED
POLICY SIZE      : 2048 K
PREFERENCE LIST: CF02      CF01
EXCLUSION LIST IS EMPTY

ACTIVE STRUCTURE
-----
ALLOCATION TIME: 02/17/96 17:02:54
CFNAME         : CF02
COUPLING FACILITY: ND02...
                  PARTITION: 0   CPCID: 00
ACTUAL SIZE     : 2048 K
STORAGE INCREMENT SIZE: 256 K
VERSION         : A8DAC970 15774B04
DISPOSITION     : DELETE
ACCESS TIME     : 0
MAX CONNECTIONS : 32
# CONNECTIONS   : 5

CONNECTION NAME  ID VERSION  SYSNAME  JOBNAME  ASID STATE
-----
DLI11            04 00040001 MVS1      DLI11    0031 ACTIVE
DLI12            05 00050001 MVS1      DLI12    0033 ACTIVE
IMS1             01 00010013 MVS1      DLI0CSA8 0036 ACTIVE
IMS2             02 00020012 MVS2      DLI0CSB8 0035 ACTIVE
IMS3             03 00030011 MVS3      DLI0CSC8 0038 ACTIVE
```

There are three parts to the output. The first part shows the status of the structure and information about the active policy. The following information from the display is especially useful:

### STRNAME

Is the name of the structure, as specified in the command.

### STATUS

Is the current status of the structure. A structure can be in a number of different states, such as ALLOCATED, NOT ALLOCATED, or ALLOCATED REBUILDING.

### POLICY SIZE

Is the size of the structure as specified in the policy.

### PREFERENCE LIST

Shows the coupling facility on which MVS allocates the structure. z/OS tries to allocate the structure on the first coupling facility listed (CF02 in the example), then the second one (CF01), and so on.

The second part of the output shows more detailed status information for the structure.

**ALLOCATION TIME**

Is the time the structure was allocated.

**CFNAME**

Is the name of the coupling facility on which the structure was allocated.

**COUPLING FACILITY**

Is the name of the coupling facility on which the structure resides.

**ACTUAL SIZE**

Is the actual size of the structure.

**DISPOSITION**

Is what the system will do with the structure in a failure situation. The IRLM structure always has a disposition of KEEP. The OSAM and VSAM structures always have a disposition of DELETE.

**MAX CONNECTIONS**

Is the maximum number of connections that can be made to the OSAM, VSAM, IRLM, shared MSGQ, or shared EMHQ structure. The maximum is 255 connections for any structure.

**# CONNECTIONS**

Is the current number of IMS systems connected to the structure.

Other information in the second part of the output is not needed for monitoring and problem determination for IMS sysplex data sharing.

The third part of the output shows which z/OS systems are connected to the structure and gives information about each use.

**CONNECTION NAME**

Is the name of a connection. z/OS will repeat the connection message text in a table to report all connections. If there are no connections, no connection table is displayed.

**ID** Is the connection identifier.

**VERSION**

Is the version number of the z/OS system that is connected.

**SYSNAME**

Is the name of the z/OS system that is connected.

**JOBNAME**

Is the name of the job associated with the connection.

**ASID**

Is the ID of the address space associated with the connection.

**STATE**

Is the status of the structure, which can be one of the following:

**FAILED PERSISTENT**

The failed persistent status occurs if the disposition of a structure is KEEP and connection to the structure is lost.

**DISCONNECTING**

IMS is in the process of disconnecting from the structure.

**FAILING**

IMS is in the process of abnormally terminating.

**ACTIVE**

IMS is connected.

**ACTIVE &**

IMS is connected, but has physically lost connectivity to the structure.

**ACTIVE OLD**

The structure is being rebuilt. IMS is connected to the old structure.

**ACTIVE &OLD**

The structure is being rebuilt. IMS is connected to the old structure, but has physically lost connectivity to it.

**ACTIVE NEW,OLD**

The structure is being rebuilt. IMS is connected to both the old and the new structure.

**ACTIVE NEW,&OLD**

The structure is being rebuilt. IMS is connected to both the old and the new structure, but has physically lost connectivity to the old structure.

**ACTIVE &NEW,OLD**

The structure is being rebuilt. IMS is connected to both the old and the new structure, but has physically lost connectivity to the new structure.

**ACTIVE &NEW,&OLD**

The structure is being rebuilt. IMS is connected to both the old and the new structure, but has physically lost connectivity to both structures.

---

## Using DBRC to control database allocation and access

DBRC allows you to control access to data by IMS systems that participate in data sharing. Using DBRC, you can modify, initiate, and delete the current status indicators in the RECON data set to change the access intent of online IMS systems and the share level of registered databases.

Your data sharing environment depends on the status of the databases and IMS systems indicated in the RECON data set.

You can modify the access intent indicator using a form of the /START or UPDATE command.

You can modify the share level indicator using a form of one of the DBRC online change commands, /RMCHANGE.

### Denial of authorization

The following occurs if DBRC (in a data sharing environment) responds to authorization requests but fails to obtain authorization for an application program:

- You receive message DFS047A identifying the database.
- IMS schedules the application program's PSB without database access. IMS abnormally terminates a BMP or MPP with abend U3303 only if it tries to access the database.
- IMS abnormally terminates batch and utility regions with abend U0047.



## Changing database access intent

Use one of the following commands to change the access intent that you declare during system definition, where *dbx* is the database name and *xx* is the new access intent:

- /START DATABASE *dbx*ACCESS=*xx*
- UPDATE DB START(ACCESS) SET(ACCTYPE(BRWS,EXCL,READ,UPD))

For the /START DATABASE *dbx*ACCESS=*xx* command, values for *xx* are:

**EX** Exclusive use  
**UP** Update access  
**RD** Read access  
**RO** Read-only access

This command is local; it affects only the IMS system on which you enter it. The GLOBAL keyword is not valid with the ACCESS= keyword. If you need to change the access level for a shared database across the sysplex, you must enter this command on each IMS system that shares the database.

The UPDATE DB START(ACCESS) SET(ACCTYPE(BRWS,EXCL,READ,UPD)) command can be either local or global.

In order to change the access intent for a DEDB, you must stop all PSBs that access any of the areas in the DEDB. You might also have to stop regions that have wait-for-input (WFI) transactions scheduled for the DEDB.

## Online DBRC commands

The /RMCHANGE command is one of a set of IMS commands that control DBRC utility functions; IMS responds to them by issuing corresponding commands directly to DBRC.

The following table shows these commands and their recommended uses.

*Table 20. DBRC commands and functions*

| Command   | Recommended use                                                                                                                                         |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| /RMCHANGE | Alter a database sharing level<br>Prevent other subsystem authorization<br>Set or remove backout-required status                                        |
| /RMDELETE | Remove database record and associated status<br>Delete subsystem with unauthorized databases<br>Remove area data set (ADS) record and associated status |
| /RMGENJCL | Generate online recovery jobs                                                                                                                           |
| /RMINIT   | Register a new database<br>Start control of an area data set                                                                                            |
| /RMLIST   | Obtain recovery and authorization information for a database or ADS                                                                                     |
| /RMNOTIFY | Reset a system status                                                                                                                                   |

If you want to change the share level of a registered database (named ORDERDB) from intraprocessor block-level data sharing (share-level 2) to interprocessor block-level data sharing (share-level 3) to allow data sharing with another IMS system in another processor, enter:

```
/RMCHANGE DBRC='DB DBD(ORDERDB) SHARELVL(3)'
```

DBRC replies with its corresponding command input (CHANGE.DB DBD(ORDERDB) SHARELVL(3)) and a series of DBRC messages.

Do not enter the /RMCHANGE, /RMDELETE, or any command that alters the status of database records in the RECON data set while a job accesses the database. Stop the database using a global /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command before modifying any RECON record for that database. Otherwise, IMS rejects the command with an error.

**Related reading:** For full descriptions of these commands and their functions, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

---

## Improving database availability

Improving database availability in a data-sharing environment involves controlling online change, reorganizing databases, and making database image copies.

### Controlling online change

If you modify a shared database online in a sharing environment, you must coordinate the sequence of /MODIFY commands across all sharing IMS systems.

1. Stop the databases you want to modify in all IMS online systems using that database.
2. Perform the online change on all sharing IMS systems
3. Restart the shared databases only when the /MODIFY command sequences are complete in all affected IMS systems.

If any of the data-sharing IMS systems also physically share any of the ACBLIB, FORMAT, or MODBLKS libraries, you must ensure that these IMS systems use the same designated active library, for example, IMS.ACBLIBA.

Each system keeps track of its active libraries in the IMS.MODSTAT data set. Performing an online change in only one IMS system could result in the individual IMS systems' using different control blocks, which can cause unpredictable errors in a data-sharing environment.

Also, if any library is shared, ensure that the Online Change Copy utility (DFSUOCU0) updates only inactive libraries.

### Reorganizing databases

While you reorganize a database (using the database reorganization utilities), you must protect it from changes by IMS systems that are authorized to use it, unless you are using HALDB Online Reorganization (OLR).

You have three ways to protect the database:

- Issue the /DBRECOVERY DB GLOBAL or UPDATE DB STOP(ACCESS) command to close the database on all IMS systems. After the database reorganization is complete, tell the MTO of each IMS system that the database is available.
- Issue the DBRC CHANGE.DB NOAUTH command to specify that the database is not authorized to participate in data sharing. After the reorganization is complete, issue the CHANGE.DB AUTH command to resume data sharing.

- Devise an MTO procedure to coordinate application activities in sharing IMS systems. This procedure would direct the MTOs to stop applications that might be active and using the database. After the reorganization is complete, the MTOs can restart applications in the appropriate online IMS systems.

The reorganization utilities communicate with DBRC. In response, DBRC does not allow further processing for the database, even after the reorganization, until you make an image copy of the database or recovery group.

The reorganization utilities might not execute if they fail to obtain authorization from DBRC for the database. DBRC denies authorization if it has authorized another IMS system to update the database (which is why you need to issue the /DBRECOVERY DB, UPDATE DB STOP(ACCESS), or CHANGE.DB commands first). If you receive message DFS044I saying that DBRC is required for the execution of the utility, and the utility abends, wait for the IMS system to terminate or, if it is an online IMS, issue a /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command.

**Related reading:** For a sample operational scenario dealing with database recovery in a database-level sharing environment, see “Reorganizing a database” on page 262.

## Making an online database image copy

When you schedule the Online Database Image Copy utility (DFSUICP0) as a BMP, DBRC manages the requirement for restrictive access. You must, however, ensure that no IMS online systems, online or batch, other than the IMS online system running the Image Copy BMP, is authorized to update the database.

For a sample operational scenario supporting image copy, see “Making an image copy of a shared database” on page 263.

---

## Maintaining normal operations

Data sharing requires additions to your normal operating procedures. The following topics present sample operational scenarios that show how to perform routine adjustments and maintenance in a database-level data-sharing environment.

**Recommendation:** Use an External Time Reference (ETR) device and set it to use Universal Coordinated Time. You can set the local time for each data-sharing IMS system separately as necessary. See *z/OS MVS Initialization and Tuning Reference* for information on using an ETR.

## Scheduling a batch updater

Given the following initial and target configurations, use this scenario to schedule a batch updater.

- Initial configuration
  - Database x (*dbx*) is registered at share level 1.
  - *dbx* is authorized for update access to an online IMS system (IMSA).
- Target configuration
  - *dbx* is authorized for update access to a batch IMS system (IMSB).
  - *dbx* is authorized for read-only access to IMSA.

Follow these steps to schedule a batch updater:

1. Prevent other IMS systems from obtaining access to *dbx* with the following command: `CHANGE.DB DBD (dbx) NOAUTH`
2. Suspend IMSA's activity on *dbx*, close *dbx*, and release IMSA's authorization with one of the following commands:

- `/DBRECOVERY DB dbxNOFEOV`
- `UPDATE DB STOP(ACCESS) OPTION(FEOV)`

You can use the `/DBRECOVERY` or `UPDATE` command for a DEDB area by specifying `AREA` instead of `DB`.

3. Allow read-only processing from IMSA with one of the following commands:
  - `/START DATABASE dbx ACCESS=RO`
  - `UPDATE DB NAME dbx START(ACCESS) SET(ACCTYPE(READ))`
4. Allow IMSB to initialize with the following command: `CHANGE.DB DBD (dbx) AUTH`
5. Start IMSB with update access to *dbx*.

## Reinstating an online updater

Given the following initial and target configurations, use this scenario to restore the initial configuration.

- Initial configuration
  - *dbx* is authorized for read-only access to an online IMS system (IMSA).
  - A batch IMS system has been terminated.
- Target configuration
  - *dbx* is registered at share level 1.
  - *dbx* is authorized for update access to IMSA.

Follow these steps to reinstate an online updater:

1. Prevent other IMS systems from obtaining access to *dbx* with the following command: `CHANGE.DB DBD (dbx) NOAUTH` When you force IMSA to release its authorization on *dbx*, another IMS system could obtain authorization to process it. The `DBRC CHANGE.DB` command with the `NOAUTH` keyword indicates (in the database record in the RECON data set) that the database is unavailable for processing.
2. Terminate IMSB.
3. Suspend IMSA's activity on *dbx*, close *dbx*, and release IMSA's authorization with one of the following commands:
  - `/DBRECOVERY DB dbx NOFEOV`
  - `UPDATE DB STOP(ACCESS) OPTION(FEOV)`
4. Allow update processing for IMSA with one of the following commands:
  - `/START DB dbx ACCESS=UP`
  - `UPDATE DB START(ACCESS)`
5. Allow access or update in IMSA with the following command: `CHANGE.DB DB dbx AUTH`

The initial configuration of the scenario is described in “Scheduling a batch updater” on page 260.

## Transferring update capability

Given the following initial and target configurations, use this scenario to transfer update capability from one sharing IMS system to another.

- Initial configuration
  - *dbx* is authorized for update access to an online IMS system (IMSA).
  - *dbx* is authorized for read-only access to another online IMS system (IMSB).
- Target configuration
  - *dbx* is authorized for read-only access to IMSA.
  - *dbx* is authorized for update access to IMSB.

Follow these steps to transfer update compatibility:

1. Prevent other IMS systems from obtaining access to *dbx* with the following command: `CHANGE.DB DBD(dbx) NOAUTH`
2. Suspend activity and close *dbx* in both IMS systems with one of the following commands:
  - `/DBRECOVERY DB dbx NOFEOV`
  - `UPDATE DB STOP(ACCESS)`
3. Allow processing in both IMS systems. Issue one of the following commands on IMSA:
  - `/START DB dbx ACCESS=RO`
  - `UPDATE DB START(ACCESS)`
4. Then issue one of the following commands on IMSB:
  - `/START DB dbx ACCESS=UP`
  - `UPDATE DB START(ACCESS)`
5. Allow access or update in both IMS systems with the following command: `CHANGE.DB DBD(dbx) AUTH` Although IMSA and IMSB now have compatible database authorizations, the database is not protected from other IMS systems until IMSA and IMSB schedule an application program with a PSB that is sensitive to the database. In order to lock out any other updaters, therefore, IMSB must be the first IMS system to obtain the authorization to process the database with update access.

## Reorganizing a database

Given the following initial and target configurations, use this scenario to reorganize a database.

- Initial configuration
  - *dbx* is registered for database-level sharing.
  - *dbx* is authorized for update access to an online IMS system (IMSA).
- Target configuration
  - *dbx* is successfully reorganized.

Follow these steps to reorganize a database:

1. Prevent other IMS systems from obtaining access with the following command: `CHANGE.DB DBD(dbx) NOAUTH`
2. Start reorganization jobs.  
Run a set of reorganization jobs, such as Scan, Unload, Reload, or Prefix Update, in a protected environment.

You must run these reorganization utilities without interference from other IMS systems. The DBRC `CHANGE.DB NOAUTH` command prohibits authorizations to the database for online and batch IMS systems, except for all reorganization utilities.

3. Authorize the Online Database Image Copy utility (DFSUICP0) initialization with the following command: `CHANGE.DB DBD(dbx) AUTH` This command allows other batch or online IMS systems to obtain access to *dbx*.

**Related reading:** For a discussion of ways to reorganize a database, see “Reorganizing databases” on page 259.

## Making an image copy with exclusive control

Given the following initial and target configurations, use this scenario to make an image copy of a database data set with exclusive control.

- Initial configuration
  - *dbx* is registered for database-level sharing.
  - *dbx* is authorized for update access to an online IMS system (IMSA).
  - *dbx* is authorized for read-only access to another online IMS system (IMSB).
- Target configuration
  - An image copy of database *dbx* (data set *y*).

Follow these steps to make an image copy with exclusive control:

1. Prevent other IMS systems from obtaining access with the following command:  
`CHANGE.DBDS DBD(dbx) DDN(y) ICON`  
Use the DBRC `CHANGE.DBDS` command with the `ICON` keyword to indicate (in the RECON data set) that an image copy is needed. This indication prevents any IMS system except the Image Copy utility from obtaining access to the database.
2. Suspend IMS system activity on *dbx* and close *dbx* in all sharing IMS systems with one of the following commands:
  - `/DBRECOVERY DB dbx NOFEOV`
  - `UPDATE DB STOP(ACCESS) OPTION(FEOV)`
3. Run the Online Database Image Copy utility (DFSUICP0).  
When this utility completes successfully, it removes the indication in the RECON data set that an image copy is needed, so that DBRC can authorize IMS systems to access the database.
4. Resume IMS system activity on *dbx* with one of the following commands:
  - `/START DB dbx`
  - `UPDATE DB START(ACCESS)`

## Making an image copy of a shared database

Because the Online Database Image Copy utility (DFSUICP0) always requests read access, it can share the database only with IMS systems that request a read-only access. If the other IMS systems need to update the database while you take an image copy, use the Concurrent Image Copy option.

Given the following initial and target configurations, use this scenario to make an image copy in a database-level sharing environment:

- Initial configuration
  - *dbx* is registered for database-level sharing.

- *dbx* is authorized for update access to an online IMS system (IMSA).
- *dbx* is authorized for read-only access to a batch IMS system (IMSB).
- Target configuration
  - The Online Database Image Copy utility (DFSUICP0) on IMSA is authorized for read access.
  - *dbx* is authorized for read access to IMSA.
  - *dbx* is authorized for read-only access to IMSB.

Follow these steps to make an image copy:

1. Prevent all IMS systems from obtaining update access with the following command: `CHANGE.DB DBD(dbx) READON`  
This command indicates (in the RECON data set) that the database cannot be accessed by IMS systems with update intent.
2. Suspend IMSA update activity on *dbx* and lower authorization to read access with one of the following commands:
  - `/DBDUMP DB dbx NOFEOV`
  - `UPDATE DB STOP(UPDATES) OPTION(FEOV)`
3. Run the Online Database Image Copy utility (DFSUICP0).
4. Resume update activity on IMSA with one of the following commands:
  - `CHANGE.DB DBD(dbx) READOFF`  
`/START DATABASE dbx ACCESS=UP`
  - `CHANGE.DB DBD(dbx) READOFF`  
`UPDATE DB NAME(dbx) START(ACCESS) SET(ACCTYPE(UPD))`

The `CHANGE.DB` command with the `READOFF` keyword indicates (in the RECON data set) that the database can be accessed by IMS systems with update intent.

## Lowering the share level of a database

Given the following initial and target configurations, use this scenario to lower the share level of a database.

- Initial configuration
  - *dbx* is authorized for update access to an online IMS system (IMSA).
  - *dbx* is authorized for read-only access to another online IMS system (IMSB).
- Target configuration
  - *dbx* is authorized for exclusive control by IMSA.

Follow these steps to lower the share level:

1. Prevent all IMS systems from obtaining access with the following command: `CHANGE.DB DBD(dbx) NOAUTH`
2. Suspend activity on *dbx* and close *dbx* in both IMS systems with one of the following commands:
  - `/DBRECOVERY DB dbx NOFEOV`
  - `UPDATE DB STOP(ACCESS) OPTION(FEOV)`
3. Allow processing in IMSA with the following command:
 

```
CHANGE.DB DBD(dbx) AUTH
/START DATABASE dbx ACCESS=EX
/START PGM
/START TRAN
```

The database is not protected from other IMS systems until IMSA schedules an application program with a PSB that is sensitive to the database. In order to



lock out all other IMS systems, therefore, IMSA must be the first IMS system to obtain the authorization to process the database with exclusive access.

IMSA accepts the change of access intent only if the target database is not authorized.

You can also achieve the same results by modifying the sharing level of the database in the RECON data set instead of changing the access intent of IMSA. Omit the /START DATABASE command with the ACCESS= keyword and perform the following step before step 3:

- Request DBRC to enforce exclusive control:

```
CHANGE.DB DBD(dbx) SHARELVL(0)
```

In this case, DBRC grants database authorization at the requested level, and IMSA owns the database exclusively.

---

## Recovering from failures

The following topics provide information to help you plan for an recover from failures in a data-sharing environment.

### Managing system logs

Recovering a shared database requires log data from each sharing IMS system. A DBRC GENJCL.RECOV command or a Database Recovery utility execution using DBRC will fail if logs need to be merged for the database.

You can use the merge function of the IMS Database Change Accumulation utility (DFSUCUM0) to merge the log data from all sharing IMS systems.

You can run the Database Change Accumulation utility (DFSUCUM0) as often as your workload requires without first establishing a complete set of log volumes. Doing so allows you to accumulate change records on a timely basis, keeping the log input to a reasonable number of volumes, without taking databases offline or bringing down IMS. Doing so can, however, increase the size of the change accumulation data set.

IMS stores records from these partial subsets of log volumes on the change accumulation data set as “spill records”. These records retain their individual update characteristics.

Before running the Database Recovery utility (DFSURDB0), you must run the Database Change Accumulation utility, specifying a complete set of log volumes for the DBDS that you need to recover. If you run the Database Change Accumulation utility with an incomplete set of logs, specify the change accumulation output data set as input to your next Database Change Accumulation run. If a change accumulation data set represents an incomplete set of logs for the DBDS that you need to recover, the data set is not valid for recovery of that DBDS.

The DBRC LIST.CAGRP command displays whether the set of log volumes for each DBDS of the change accumulation group is complete or incomplete.

You can run DFSUCUM0 with a valid set of log volumes at any time in order to minimize the size of the change accumulation data set. To establish a valid set of logs in a data-sharing complex, use the following procedure:

1. Suspending operation of all batch IMS systems



2. Using the /DBRECOVERY DB or UPDATE DB STOP(ACCESS) command to terminate database access for all but one online IMS system
3. Forcing a log switch on the remaining online IMS system

## Planning recovery procedures

Recovery in a data-sharing environment is similar to standard IMS recovery.

Both standard IMS recovery and recovery in a data-sharing environment involves these primary tasks:

- Setting up the mechanisms of recovery (logging, taking checkpoints, keeping records)
- Setting up the operational procedures to be followed in the event of situations that require recovery

The logging and checkpoint mechanisms of online IMS systems in a nonsharing environment are also active in a data-sharing environment. These include:

- System log data sets and the use of the WADS and restart data sets
- Application program, system, and message-queue checkpoints
- Making image copies of databases

The primary difference between nonsharing and data-sharing environments is in their degree of reliance on DBRC. DBRC helps control the data-sharing environment; it does not merely keep records.

As in a single IMS system, a failure in a data-sharing environment might require offline (utility-type) recovery as well as restart. After a failure in a data-sharing environment, however, you must consider and perform recovery actions not only in the failing system, but also in other sharing systems. In addition, each IRLM, DBRC, and coupling facility structure in the environment introduces another potential point of failure and another component that you might need to include in a recovery and restart procedure.

Recovery for the message queue data sets, (in a nonshared-queues environment) log data sets, and system data sets are no different in a data-sharing environment:

- To recover the log, use the Log Recovery utility (DFSULTR0).
- To recover the message queue data set, request at IMS restart that IMS rebuild the queues. You need to supply the system log data sets since the last system checkpoint that recorded the queue contents (SNAPQ, DUMPQ, or PURGE).
- To recover a system data set other than the message queue data set, update the latest image copy. You must recover the RECON data set for data sharing to continue.

## Recovery facilities

Dynamic backout, batch backout, and forward recovery facilities are different in data-sharing and nonsharing environments.

### Dynamic backout

As in a nonsharing environment, an online IMS system in a data-sharing environment dynamically backs out an application program's uncommitted database changes and discards its uncommitted output messages if the application program fails or requests backout with a rollback call. In a data-sharing environment, however, IRLM locks and RECON status indicators ensure integrity by protecting uncommitted changes from sharing IMS systems.

After an application program failure, operation of the system and other application programs continues uninterrupted.

If the application program is a batch updater in a block-level data-sharing environment, IMS dynamically backs out changes in any of the following situations:

- When the application program issues a rollback call (ROLL, ROLB, or ROLS).
- When the batch region uses system logging allocated to DASD, you specify BKO=Y as an execution parameter, and the application program issues a ROLB or ROLS call or a DL/I-detected error occurs.
- When some types of resource shortages and deadlocks involving two or more batch update application programs occur, IMS performs dynamic backout for one or more of the batch application programs; all continue executing, provided that the involved regions use DASD logging and you specify BKO=Y.

In these situations, you are not involved in recovery actions.

For either online or batch IMS systems, if the database is shared at the database level, other application programs might read uncommitted data before the database is dynamically backed out. Backout affects only the updating application program's data in the database—not the output of any other application program that used this uncommitted data. This is a consequence of the nature of PROCOPT=GO processing, rather than of the nature of database backout in a data-sharing environment.

If a dynamic backout for a database fails, you must back out the database using batch backout. Then, you must restart that database using a /START DB or UPDATE DB START(ACCESS) command.

### **Batch backout**

If the failing application program is a batch updater involved in block-level data sharing, IMS does not always automatically back out its uncommitted database changes.

Just as in a nonsharing environment, you must recover the database if any of the following is true:

- The batch region does not use DASD logging
- You do not specify BKO=Y as an execution parameter
- The failure is not a DL/I-detected error

In a block-level-data-sharing environment, however, you must back out database changes of all batch update application programs using the Batch Backout utility (DFSBB00).

During block-level sharing, if an application program deadlock occurs involving only batch update jobs, IMS abnormally terminates one of these jobs with abend U0777. If IMS can dynamically backout changes, other application programs continue processing. Otherwise, IMS abnormally terminates the other application programs with abend U3303.

You must recover all of the batch jobs that did not have dynamic backout by executing the Batch Backout utility for each of them. Because the locks held by these batch jobs could affect access of other batch jobs or online IMS systems, you should begin batch backouts promptly.

**Restriction:** For recovery of batch IMS systems that participate in block-level data sharing:

- The Batch Backout utility requires an IRLM for its execution.
- The Batch Backout utility backs out to the last successful checkpoint.

### **Forward recovery**

Recovery of a database in a data-sharing environment is similar to recovery in a nonsharing environment. In both environments, use the Database Recovery utility (DFSURDB0) and give it the most recent image copy of the database, along with all pertinent log data sets used since you made that image copy.

Block-level data sharing, however, can require one additional step. Because more than one IMS system might update the database, you need the log from each IMS system to reconstruct the database. Furthermore, because these updates are most likely to be concurrent, you cannot submit those logs sequentially to the Database Recovery utility. Instead, you must first merge them by running the Database Change Accumulation utility (DFSUCUM0). Be sure to process all pertinent logs.

Even if you run change accumulation regularly, you must run the utility again before recovering the database so you merge the most recent logs with earlier logs.

However, you do not need to establish a valid set of logs for the Database Change Accumulation utility (DFSUCUM0). You can run DFSUCUM0 at any time without suspending operation of batch IMS systems, terminating database access for all but one online IMS system, or forcing a log switch on the remaining online IMS system. See “Managing system logs” on page 265.

## **Recovery without DBRC**

If you perform any recovery-related actions when DBRC is not running, such as making database image copies, problems could arise because DBRC would be unaware of changes in status. You must, therefore, specifically inform DBRC of such changes.

DBRC offers several commands for this purpose; see “Online DBRC commands” on page 258.

Because restart might be required, you should, for example, use the Utility Control Facility (UCF) without DBRC being active. But, if you use UCF for any recovery-related work, you must specifically notify DBRC of status changes afterward.

## **Restart after IMS failure**

Restart an IMS system in a data-sharing environment in the same way as in a nonsharing environment.

- Restart IMS normally, using the /NRESTART command, if you were able to shut the IMS system down normally.
- Emergency restart IMS, using the /ERESTART command, if the IMS system failed.

In a data-sharing environment, however, consider that if the associated IRLM also stops or fails, you must restart that IRLM before you start IMS in a block-level data-sharing environment. See “Starting the IRLM” on page 249.

## Restart after DBRC failure

Because DBRC runs under the control of IMS, if IMS abnormally terminates, IMS tries to release DBRC buffers, and to close the system log.

After you correct the DBRC problem, restart IMS using the `/ERESTART` command.

## Recovery involving IRLM

Failure in one part of a data-sharing configuration does not usually cause the other parts to fail.

For example:

- If one z/OS system fails because of a hardware problem or an abnormal termination of the z/OS operating system, the other systems—with their IRLMs and other IMS systems—continue to run. In this case, the surviving IRLMs are notified of the failure by z/OS and then provide the necessary retained lock protection for locks owned by the failing member.
- If an IRLM fails, its associated IMS online systems continue to run while all active application programs receive U3303 abends, and IMS suspends transaction scheduling, allowing backouts to complete. Batch IMS systems terminate to protect database integrity. The surviving IRLMs are notified of the failure by z/OS and then provide the necessary retained lock protection for locks owned by the IMS on the failing IRLM.

If an IRLM fails in an IMS DBCTL environment, in-doubt threads receive a U3303 abend on their next DL/I calls. The in-doubt threads remain active until phase 2 of sync-point processing. After processing phase 2 commit or abort, the thread is terminated. The IMS DBCTL system tries to reconnect to IRLM after all in-doubt threads are terminated.

- If an IMS system fails, its associated IRLM and other IMS systems continue to run, with locks owned by the failing IMS protected by retained locks created by its IRLM.
- If DBRC fails, its IMS system abnormally terminates, but the associated IRLM and the other IMS systems and their IRLMs continue to run.

## Transaction recovery in online IMS systems

IMS automatically backs out transactions if the IMS online system determines that transaction processing cannot proceed because of problems directly or indirectly to do with data sharing, or because the resource is unavailable, and if you specified `SERIAL=NO` in the `TRANSACT` system definition macro (or accepted it as the default).

If you specify `SERIAL=YES` in the `TRANSACT` system definition macro, IMS puts the transaction back on the front of the transaction queue to be processed first-in-first-out, and stops the transaction (USTOP state).

If the transactions do not use the `DL/I INIT STATUS GROUPA` call, IMS places transactions on either the suspend queue or the transaction queue when:

- A transaction requests a lock held by an IMS system identified to a failed IRLM.
- A transaction requests a lock held by a failed IMS system.
- The transaction attempts to access a database at a level higher than that for which the IMS system is authorized.
- A transaction requires a database that is stopped or locked.
- The MTO enters a `/DBRECOVERY DB` or `UPDATE DB STOP(ACCESS)` command.

IMS sends message DFS3324I, saying that the transaction is suspended, for each individual transaction that is suspended.

IMS saves the suspend queue across any restart, except a cold start, and all transactions on the suspend queue are eligible for reprocessing. When IMS reprocess the transactions, IMS releases them to the message queues and schedules them in the normal way.

Several conditions prompt IMS to attempt to reprocess these transactions:

- A restart of the online IMS system, including after an XRF takeover
- A communication from another IMS system performing restart or from a batch IMS system that has completed backout
- A reconnection of an IMS online system to an IRLM
- The MTO command, /DEQUEUE SUSPEND
- The MTO commands, /START DB or UPDATE DB START(ACCESS) to start the database

IMS issues message DFS3336I when it attempts to reprocess suspended transactions.

### **Restart with a single failing IRLM**

If the only IRLM in a data-sharing configuration fails: the associated IMS online systems suspend all transaction scheduling, IMS performs backout for affected transactions, IMS places the transactions on a suspend queue or transaction queue, depending on whether you specified SERIAL=NO or SERIAL=YES in the TRANSACT macro and IMS stops shared databases.

IMS issues DFS2012I messages for each database that has been stopped automatically because of the failure. IMS abnormally terminates batch IMS systems associated with the IRLM. Batch database backout is required for only those batch executions that perform database updates and do not use dynamic backout.

After you restart the failed IRLM, enter a z/OS MODIFY command, specifying the RECONNECT keyword, for each IMS online system. This command reestablishes the data-sharing environment, and initiates reprocessing for the suspend queue.

**Related reading:** For more information about transaction recovery, see “Transaction recovery in online IMS systems” on page 269)

### **Restart in a sysplex data-sharing environment**

In a sysplex data-sharing environment, an IRLM is informed of z/OS failure, IMS failure, failure of another IRLM, or loss of connection to physical locking resources.

The IRLM affected by the failure disconnects from the data-sharing group. The remaining IRLMs in the data-sharing group are referred to as surviving IRLMs.

When an IRLM is forced to leave a data-sharing group for any reason, that IRLM issues the DXR136I message, which says that an IRLM has disconnected from the data-sharing group.

The surviving IRLMs issue the DXR137I message, which says the group status has changed, and that an IRLM has been disconnected from the data-sharing group.

No operator action is required for either message. As in a nonsysplex data-sharing environment, IRLM protects resources that are locked from being accessed by other IMS systems.

Normal restart procedures can be used to restart the IRLM and its associated IMS.

---

## Data sharing in a sysplex environment

For the most part, nonsysplex data sharing concepts apply equally to sysplex data sharing. There are exceptions, which are discussed in these topics.

These topics assume you are familiar with z/OS sysplex concepts and terminology.

### Sysplex data-sharing concepts and terminology

When multiple IMS systems share data across more than two z/OS images, it is called *sysplex data sharing*.

Each z/OS image involved in sysplex data sharing must have at least one IRLM and a corresponding Fast DB Recovery (FDBR) region. Each IRLM is connected to a coupling facility. The coupling facility, as explained in “Coupling facility” on page 273, is used to maintain data integrity across IMS systems that share data.

#### Buffer invalidation

When a database block or CI is updated by one IMS, IMS must make sure all copies of this block or CI in other IMS buffer pools are marked invalid.

To do this, the updating IMS issues an invalidate call. This call notifies the coupling facility that the content of a buffer has been updated and that all other copies of the buffer must be invalidated. The coupling facility then invalidates the buffer for each IMS that registered an interest in the buffer. This process is called *buffer invalidation*. Whenever IMS tries to access a buffer, it checks to be sure that the buffer is still valid. If the buffer is invalid, IMS rereads the data from DASD. Data integrity is thereby maintained.

Buffer invalidation works in all IMSplex database environments: DB/DC, DBCTL, and DB batch. In the sysplex environment, IMS supports buffer pools for VSAM, VSAM hiperspace, OSAM, and OSAM sequential buffering buffers.

#### Data-sharing groups

As with non-sysplex data sharing, the concept of a data-sharing group applies. The following figure shows a sample data-sharing group.




Figure 35. Sample data-sharing group

The data-sharing group has two operating systems. Each operating system contains an IMS control region, an FDBR, and an IRLM. The IMS systems in each operating system share the following resources:

- Databases
- A RECON dual copy data set
- One or more coupling facilities
- A single IRLM lock table structure (lock structure) in the coupling facility
- OSAM and VSAM buffer invalidate structures in the coupling facility (hereafter called an OSAM or VSAM structure)



Fast database recovery regions are not shared.

Communication in the data-sharing group is through the IRLMs connected to the coupling facility. Without a coupling facility, no valid sysplex data-sharing environment exists; IRLM does not grant global locks, even for non-sysplex data sharing. IRLM requires a coupling facility for sysplex data sharing. If a coupling facility is not available, IRLM allows only non-sysplex data sharing.

IMS connects to the structures in the coupling facility. Sysplex data sharing uses the OSAM and VSAM structures for buffer invalidation. This differs from non-sysplex data sharing, which does buffer invalidation using broadcasts (notifies) to each IMS. Sysplex data sharing uses the lock structure in the coupling facility to establish and control the data-sharing environment.

A data-sharing group is defined using the CFNAMES control statement in the DFSVSMxx member of the IMS PROCLIB data set. The lock, OSAM, and VSAM structures are named on the CFNAMES control statement.

## **Coupling facility**

Figure 35 on page 272 shows the coupling facility and the three structures in it that are used for sysplex data sharing.

Although the figure shows a single coupling facility, more than one are possible. Additional coupling facilities can be defined for backup. Or, to increase throughput, structures can be split across coupling facilities; the lock structure, for example, can be put on one coupling facility and OSAM and VSAM structures on another.

The OSAM and VSAM structures are used for buffer invalidation. For each block of shared data read by any IMS system connected to the coupling facility, an entry is made in the OSAM or VSAM structure. Each entry consists of a field for the buffer ID (known to z/OS as the resource name) and 32 slots. The slots are for IMS systems to register their interest in an entry's buffer. This makes it possible for as many as 32 IMS systems to share data. Note that the sysplex data-sharing limit of 32 IMS systems is the number of IMS systems that can connect to a structure; it is not the number of IMS systems that are running. It is also not the number of total connections that can be made to the coupling facility which is 255.

The lock structure is used to establish the data-sharing environment. For a data-sharing group, the first IMS system to connect to a lock structure determines the data-sharing environment for any IMS system that subsequently connects to the same lock structure. When identifying to the IRLM, IMS passes the names of the coupling facility structures specified on the CFNAMES control statement, plus the DBRC RECON initialization time stamp (RIT) from the RECON header. The identify operation fails for any IMS system not specifying the identical structure names and RIT as the first IMS.

If a structure fails (or you initiate a manual rebuilding of the structure), IMS tries to rebuild it. If IMS is unsuccessful or if the connection to a structure is lost, IMS quiesces data-sharing transactions and stops data sharing. After the rebuilding is complete, IMS tries to reconnect to the structure. If IMS reconnects successfully, IMS continues processing data-sharing transactions. If IMS fails to reconnect, data sharing remains stopped and IMS waits to reconnect until it is again notified that coupling facility resources are available.



**Note:** If a batch job is connected to a structure when IMS attempts to rebuild the structure, IMS terminates the batch job and issues an abend U3303.

**Related reference:**

 zSeries Processor Resource/Systems Manager Planning Guide

 z/OS: Defining a coupling facility

## XRF and sysplex data sharing

XRF can be used in a sysplex data-sharing environment only if the coupling facility is available and connected when the alternate system is started.

**Restriction:** You cannot use Fast DB Recovery for XRF subsystems.

## Sample sysplex data-sharing configurations

The following figures show several possible configurations for sysplex data sharing.

The DBRC RECON data set and the IMS database are not shown in these figures but, as in Figure 35 on page 272, RECON and the database are shared by the IMS systems in the sysplex data-sharing group.

### Three structures on one coupling facility

The following configuration shows multiple z/OS images running in the sysplex data-sharing environment using IRLM. IMS is running on each z/OS image. One coupling facility is being used. The lock, OSAM, and VSAM structures are on the same coupling facility.




Figure 36. Sysplex data-sharing configuration with three structures on one coupling facility

### Three structures on two coupling facilities

The following configuration illustrates multiple z/OS images running in the IMSplex using IRLM. IMS is running on each z/OS image. Two coupling facilities are being used. The lock structure is on coupling facility 1. The OSAM and VSAM structures are on coupling facility 2.




Figure 37. Sysplex data-sharing configuration with three structures on two coupling facilities

### Three structures on one coupling facility with backup

The following configuration shows multiple z/OS images running in the IMSplex using IRLM. IMS is running on each z/OS image. Two coupling facilities are being used. The lock, OSAM, and VSAM structures are on coupling facility 1. Coupling facility 2 is used as a backup so that if coupling facility 1 fails, structures can be rebuilt on coupling facility 2.




Figure 38. Sysplex data-sharing configuration with three structures on one coupling facility with backup

### Three structures on one coupling facility with backup and XRF

The following configuration shows six z/OS images running in the IMSplex using IRLM. Three active (XRF) IMS systems are running, and each active IMS has an alternate system. The lock, OSAM, and VSAM structures are on the coupling facility.




Figure 39. Sysplex data-sharing configuration with three structures on one coupling facility with backup in an XRF environment

### One structure on one coupling facility

The following configuration shows multiple z/OS images running in a data-sharing IMSplex using IRLM. IMS is running on each z/OS image. One coupling facility is being used. The lock structure is on the coupling facility. No OSAM or VSAM structures are on the coupling facility. This configuration results in a data-sharing IMSplex using the notify protocol for buffer invalidation.




Figure 40. One structure on one coupling facility resulting in sysplex data sharing

## When to use sysplex data sharing

The two primary reasons for moving to a sysplex data-sharing environment involve cost and availability. With z/OS sysplex data sharing, you can take advantage of the lower cost of the CMOS-based hardware, spreading the workload to multiple CPUs while maintaining one logical view of your databases; this can be done without making any application software changes.

Availability in this environment is improved, because if one CPU in a sysplex is lost, the workload can be moved to the remaining CPUs.

## Modifying batch jobs for sysplex data sharing

Each IMS connected to an OSAM or VSAM structure uses one connection or slot to each coupling facility structure. Therefore, every batch job running in data-sharing mode uses a slot.

Any jobs in the same data-sharing group (started after the combination of active online and batch jobs) that exceed the maximum are unable to run. If you convert these batch jobs to batch-oriented BMP jobs, jobs are more likely to be able to run because BMP jobs run under the online region coupling facility connection.

For more information about converting batch jobs to batch-oriented BMP jobs, see *IMS Version 11 Application Programming*.

## Calculating the size of coupling facility structures

The size of structures in the coupling facility is determined using a z/OS formula. You must calculate certain IMS values in this formula.

**Recommendation:** Use transaction rates of peak processing periods for calculating structure sizes. Doing so allows you to avoid resetting structure sizes. Overestimating sizes causes no problems, while underestimating sizes can cause abends.

### **Calculating the size of the lock structure**

The size of a coupling facility lock structure is dependent on a variety of factors that are unique to your installation, such as transaction rates, update rates, contention rates, commit frequency, and so forth. Consequently, there is no way to calculate the optimum size of a lock structure that does not involve some estimation, monitoring, and adjustment.

The simplest method is to start with a likely size that is a power of 2, such as 256 MB or 512 MB. Then monitor your update and contention rates. Tune the lock structure size to lower contention to an acceptable level and to ensure that enough room is available for modify (update) locks. Lock structures that are too big can also adversely affect performance.

You can also use the z/OS Coupling Facility Structure Sizer Tool (CFSizer). CFSizer is a web-based application that calculates the structure size based on the input data you provide. If you use CFSizer, you still need to monitor your update and contention rates and adjust the lock structure size as necessary. To use the CFSizer tool, go to <http://www.ibm.com/systems/support/z/cfsizer/>.

After setting the initial value, monitor the use of the lock structure using IRLM messages, some of which suggest that you increase the size of the structure. You will begin to receive IRLM messages when 50% of the current lock structure is used.

The coupling facility lock structure contains two parts. The first part is a lock entry table used to determine if there is inter-IMS read/write interest on a particular hash class (resources that hash to a particular place in the lock table). The second part is a list of the update locks that are currently held (sometimes called a modify lock list or record list table). The division of the lock structure storage between these two components can be controlled by the user through the IRLM DXRJPROC procedure or by an IRLM MODIFY command. If the user does not specify how the structure is to be split, then IRLM will attempt to divide it with a 1:1 ratio between lock table entries and record list storage. The total size of the lock structure must be large enough to prevent performance problems by limiting hash contention and to prevent failures resulting from lack of record list storage to write a MODIFY entry (RLE). Proper specification for the number of lock table entries can help avoid hash contention.

IRLM reserves 10% of the record table entries for “must complete” functions (such as rollback or commit processing) so that a shortage of storage does not cause an IMS system failure. However, if storage runs short in the record table, there can be an impact on availability (transactions are terminated), response time, and throughput.

### **Specifying the lock entry size**

The field LOCK ENTRY SIZE of the installation panel DSNTIPJ determines the amount of space required for lock contention control information (that is, individual entries in the lock table). The lock entry size and the number of lock table entries of the first IRLM to join the group causing structure allocation determine the storage size needed for the lock table and the lock table entry width.

for the whole group. The default is two bytes, which is probably the number you want, unless you immediately create a data sharing group of seven or more members. By restricting each lock entry to two bytes, you maximize the amount of LTE space available from the define structure size. This can help avoid false contention.

### **Storage estimate for the lock structure**

For installation planning purposes, the initial size of the lock structure is based on how much updating you do.

**Recommendation:** If you do not specify the `LTE=` keyword in the IRLM DXRJPROC procedure, choose a value for the `INITSIZE` that is a power of 2. This enables IRLM to allocate the coupling facility storage so that half will be used for lock table entries and the remainder for the record table entries. If a 1:1 split occurs and total size is not a power of 2, you may experience severe shortage of space for the record table entries, resulting in IMS or possible IRLM failures. (This will occur because the number of lock table entries requested on `CONNECT` must be a power of 2.) The record table is susceptible to storage shortages if the structure is too small or if the allocation of the lock table leaves too little storage for the record table.

When specifying a value for the `LTE=` parameter in the IRLM DXRJPROC procedure, or when issuing the `IRLM MODIFY SET,LTE=` command, you should monitor XES contention rates to determine the optimum value for your normal environment. If the contention rates appear too high, then increase the `LTE=` value to the next power of 2, keeping in mind that any increase in the size of the lock table will cause a corresponding decrease in the record table, unless the structure size is also increased. If you have little contention and want more storage available for record table entries, then decrease the `LTE=` value by a power of two. Anytime the number of lock table entries are decreased, it is good to monitor contention rates for a period of time.

Since the structure allocation is performed at `CONNECT` time, any change made to the `LTE=` value does not take effect unless the group is terminated, structure forced and the group restarted or a `REBUILD` is done. Also, the `LTE=` value of the first IRLM to `CONNECT` dictates the coupling facility attributes used by the group.

### **Automatic altering of the size of the lock structure**

z/OS can automatically expand or contract the size of a lock structure in the coupling facility if it needs storage space. The `ALLOWAUTOALT` parameter in the CRFM policy specifies whether system-initiated alters (automatic alters) are allowed for the lock structure. If no competing resource is using the coupling facility, `ALLOWAUTOALT(YES)` might be desirable to allow the lock structure to change as the workload grows. However, with the parameter set to `YES`, another resource can force the lock structure to decrease, possibly causing an IRLM out-of-storage condition in the lock table during coupling facility use. In this case, the application terminates abnormally with a U3307 abend. In addition, the IMS control region might terminate with a U0113 or U1027 abend.

If you need the lock structure to be stable, use `ALLOWAUTOALT(NO)`.

### **Calculating the size of OSAM and VSAM structures**

For sysplex data sharing, the size of the structure required by each access method depends on the number of buffers defined to the access method by each IMS in the



data-sharing group. The buffers for all IMS control regions and batch jobs that are data sharing and registered to IRLM must be counted.

The following formula is for calculating the count of OSAM buffers:

OSAM buffer count = #osambfrs/IMS1 + #osambfrs/IMS2 + ... # osambfrs/IMSn

The following formula is for calculating the count of VSAM buffers:

VSAM buffer count = #vsambfrs/IMS1+ #vsambfrs/IMS2 + ... #vsambfrs/IMSn

Here is an example of sizing a VSAM structure with 2 IMS systems (IMS1 and IMS2):

- IMS1 has a total of 700 buffers (400 VSAM buffers plus 300 Hiperspace™ buffers) based on the following specifications:

VSRBF=512,30  
VSRBF=1024,20,I  
VSRBF=1024,10,D  
VSRBF=2048,40  
VSRBF=4096,100,HS100,HSR  
VSRBF=8192,100,HS100,HSR  
VSRBF=16384,50,HS50,HSR  
VSRBF=32768,50,HS50,HSR

- IMS2 has a total of 400 buffers based on the following specifications:

VSRBF=512,30  
VSRBF=1024,20,I  
VSRBF=1024,10,D  
VSRBF=2048,40  
VSRBF=4096,100  
VSRBF=8192,100  
VSRBF=16384,50  
VSRBF=32768,50

By adding the total number of buffers for both IMS systems (700 + 400), you get the number to be used for sizing the VSAM structure; 1100.

**Recommendation:** Review the size of the OSAM and VSAM Structures any time OSAM or VSAM buffer pool definitions are changed. You might encounter poor data sharing performance if these structures are not large enough.

An OSAM structure may optionally be used for caching data. If data caching is not used, the structure contains only directory entries. If data caching is used, the structure contains directory entries and data elements. In this case the size of the structure must allow for data elements. OSAM data is stored in the structure as multiples of 2 KB data elements. Using data caching requires a directory-to-element ratio to be specified. This ratio causes the structure to be subdivided with directory entries and data elements. The ratio is specified with the CFOSAM= keyword on the CFNAMES parameter statement.

For more information on the ratio is specified with the CFOSAM= keyword on the CFNAMES parameter statement, see *IMS Version 11 System Definition*.

**Requirement:** The OSAM buffer count must include any sequential-buffering buffers that are defined. The VSAM buffer count must include any hiperspace buffers that are defined.

## **z/OS formula**

After you have calculated the OSAM buffer count, put the result in the TDEC field of the z/OS cache structure size formula. Then do the same for the VSAM buffer count. (The calculation is done separately for each access method.) TDEC is the count of IMS buffers (either OSAM or VSAM).

For information on OSAM database coupling facility caching, see *IMS Version 11 System Definition*.

## **Changing the size of OSAM and VSAM structures**

You can dynamically change the size of OSAM and VSAM cache structures or reapportion objects within a structure by using Structure Alter.

You can also automatically perform Structure Alter when a structure fills to change the structure size and element-to-entry ratio. A coupling facility level of 1 is required to support Structure Alter.

For more information about Structure Alter, see “Using structure alter for CQS” on page 179.

## **Recovering from sysplex data-sharing failures**

Multiple failures can occur in a sysplex data sharing environment.

If multiple failures occur simultaneously (for example, IRLM and OSAM structures fail on the same coupling facility), each component (IRLM, IMS, and so on) takes the appropriate recovery action.

### **Coupling facility connection failures**

Various types of problems can cause the connection between IMS and the coupling facility to be lost. For example, the connection is lost if the physical link (a fiber optic cable) is cut, after power failures, hardware check stops, coupling facility microcode failures, and similar events.

If there are multiple links between an IMS and the coupling facility, loss of one link has no effect on IMS or the IRLM. However, if the link that is lost is the last link, it must be recovered before data sharing on that IMS can be restarted. Processing still continues for all other IMS systems in the data-sharing group.

When the last connection between IMS and the coupling facility is lost, what happens depends on the IMS operating environment:

#### **Batch Environment**

All IMS batch jobs that were connected to the coupling facility terminate with a U3303 abend.

#### **DB/DC or DBCTL Environment**

If the connection was between an IMS and an OSAM or VSAM structure, IMS drops out of data sharing. Any transactions trying to access shared data terminate with a U3303 abend. However, transactions accessing nonshared databases can still be processed.

If the connection that failed was between an IRLM and the IRLM lock structure, the IRLM disconnects from the data-sharing group. Any transactions trying to access shared data terminate with a U3303 abend.

When connection between the coupling facility and IMS is reestablished, z/OS notifies IMS and IMS resumes data sharing.

The following figure shows the connection loss between IMS, OSAM, and VSAM structures as a result of a coupling facility connection failure.




Figure 41. Connection loss between IMS and OSAM and VSAM structures

The connection that is lost is the link between the MVS2 and coupling facility CF02. When the last link fails:

- IMS2 drops out of data sharing and only nonshared databases can be accessed on the MVS2 system. IMS1 and IMS3 continue data sharing.
- Transactions trying to access shared data on MVS2 terminate with a U3303 abend.

In this circumstance, you receive the DATA SHARING STOPPED message. You can enter the z/OS DISPLAY XCF,STRUCTURE,STRNAME= command to determine what failed. The displayed output, as shown in the following example, shows that the MVS2 link to the coupling facility failed. MVS1 and MVS3 links are still active.

```
11:44:37.84 D XCF,STRUCTURE,STRNAME=OSAMESXI
```

```
⋮
```

```
& AMPERSAND DENOTES CONNECTOR WHO LOST CONNECTIVITY TO STRUCTURE
CONNECTION NAME  ID  VERSION  SYSNAME  JOBNAME  ASID  STATE
-----
```

|       |                  |                       |
|-------|------------------|-----------------------|
| IMS1  | 03 0003000D MVS1 | DLIOCSA8 0066 ACTIVE  |
| &IMS2 | 01 0001001D MVS2 | DLIOCSB8 0066 FAILING |
| IMS3  | 02 00020014 MVS3 | DLIOCS8 0066 ACTIVE   |

In the online environment, the MVS2 link to the OSAM structure is reestablished, MVS notifies IMS, and data sharing resumes. In the batch environment, you need to resubmit the batch job when the link is reestablished. You will receive the CF INITIALIZATION COMPLETE message, indicating IMS has resumed data sharing. If you issue a DISPLAY XCF,STRUCTURE command, display output shows that MVS2 is connected (ACTIVE).

## Coupling facility structure failures

When a coupling facility structure fails in a IMS batch environment, all batch jobs terminate. In a DB/DC or DBCTL environment, all IMS systems in the data-sharing group temporarily suspend data sharing.

When a structure fails, what happens depends on the operating environment:

### Batch environment

All IMS batch jobs that were connected to the coupling facility terminate with a U3303 abend.

### DB/DC or DBCTL environment

All IMS systems in the data-sharing group temporarily suspend data sharing. If the structure that failed was an OSAM or VSAM structure, IMS tries to rebuild the structure. If the rebuild is successful, all IMS systems connect to the new structure and data sharing continues. If the structure that failed was an IRLM structure, all IRLMs try to rebuild the IRLM structure. If the rebuild is successful, all IRLMs connect to the new structure and locking for data sharing continues.

**Example:** Figure 41 on page 284 shows the configuration for an OSAM structure failure. When the OSAM structure on coupling facility CF02 fails, the following happens:

- All IMS systems suspend data sharing. If you display the status of the structure at this point (using the DISPLAY XCF,STRUCTURE,STRNAME=OSAMESXI command), output shows that the structure is in the quiesce phase of rebuild, as shown in the following example.

```
08:50:59.29 D XCF,STRUCTURE,STRNAME=OSAMESXI
08:51:13.32 IXC360I 08.50.59 DISPLAY XCF 143
STRNAME: OSAMESXI
STATUS: REASON SPECIFIED WITH REBUILD START:
        STRUCTURE FAILURE
        REBUILD PHASE: QUIESCE
POLICY SIZE      : 2048 K
PREFERENCE LIST: CF02      CF01
EXCLUSION LIST IS EMPTY
```

⋮

```
* ASTERISK DENOTES CONNECTOR WITH OUTSTANDING REBUILD RESPONSE
CONNECTION NAME  ID VERSION  SYSNAME  JOBNAME  ASID STATE
-----
*IMS1            03 00030015 MVS1      DLIOCSA8 003A ACTIVE
*IMS2            01 00010026 MVS2      DLIOCSB8 0038 ACTIVE
*IMS3            02 0002001D MVS3      DLIOCS8 0038 ACTIVE
```

- When all IMS systems suspend data sharing, a new structure is built. All IMS systems are then connected to the new structure and data sharing resumes. If you display the status of the new structure after rebuilding, output is similar to that shown in the following example.

```
08:52:39.78 D XCF,STRUCTURE,STRNAME=OSAMSESXI
08:52:40.43 IXC360I 08.52.39 DISPLAY XCF 168
          STRNAME: OSAMSESXI
          STATUS: ALLOCATED
          POLICY SIZE      : 2048 K
          PREFERENCE LIST: CF02      CF01
          EXCLUSION LIST IS EMPTY

ACTIVE STRUCTURE
-----
ALLOCATION TIME: 03/25/94 08:51:27
CFNAME          : CF01
COUPLING FACILITY: ND01...
                  PARTITION: 0      CPCID: 00
ACTUAL SIZE     : 2048 K

:
```

| CONNECTION NAME | ID | VERSION  | SYSNAME | JOBNAME  | ASID | STATE  |
|-----------------|----|----------|---------|----------|------|--------|
| IMS1            | 03 | 00030015 | MVS1    | DLIOCSA8 | 003A | ACTIVE |
| IMS2            | 01 | 00010026 | MVS2    | DLIOCSB8 | 0038 | ACTIVE |
| IMS3            | 02 | 0002001D | MVS3    | DLIOCS8C | 0038 | ACTIVE |

## Failure in rebuilding a coupling facility

You can initiate a rebuild of the coupling facility by issuing the z/OS SETXCF START,REBUILD command.

z/OS can also initiate a rebuild in response to:

- A coupling facility structure failure
- A connection failure if the active System Failure Management (SFM) policy threshold for rebuilding the coupling facility has been reached

If the rebuild fails, IMS stops data sharing. Transactions trying to access shared data terminate with a U3303 abend. If the rebuild failure was associated with an OSAM or VSAM structure, transactions accessing nonshared databases can still be processed. If the failure was associated with an IRLM structure, the IRLM disconnects from the data-sharing group.

## Coupling facility connection failure during IMS restart

If the connection of the active system to the coupling facility fails during IMS restart, restart processing continues.

If the failure is on the XRF alternate subsystem and the alternate subsystem is unable to connect to the coupling facility, restart cannot proceed.

## Failure when IMS tries to identify to an IRLM

A failure occurs if IMS cannot identify itself to the IRLM. When this happens, IMS issues a message indicating the reason for the failure and then abends.

You can gather information to determine what action to take using the DISPLAY XCF,STRUCTURE and MODIFY irlmprocx,STATUS commands.

Except for this identify failure, IRLM failure processing is the same as with nonsysplex data sharing. IMS:

- Backs out all uncommitted database changes.
- Removes authorization of all shared databases through DBRC. The system is stopped.
- Waits for the IRLM to be restarted.

### **z/OS and other failures requiring system recovery**

If z/OS fails or other failures occur that require system recovery (for example, a power outage), IMS processing is the same as it is in the nonsysplex data-sharing environment.

For the failing system, operations consist of:

- Restarting z/OS
- Restarting the IRLM
- In the batch environment, backing out batch jobs if necessary
- In the IMS DB/DC or DBCTL environment, emergency restarting IMS (using the /ERESTART command)

You can restart z/OS and the IRLM on a different system. Also, batch backout and emergency restart can be done on any z/OS and IRLM in the data-sharing group.

For the non-failing systems:

- Locks are retained for the failed IMS.
- Locks are released as the failed IMS restarts and recovers its databases.
- Database buffers for the failed IMS are invalidated.
- Database error conditions are sent to the other IRLMs.

### **Synchronizing IMS restart after a multiple IMS system failure**

Multiple IMS system failures can occur in a data-sharing environment. When this happens, you might want to synchronize the startup of all of the failed IMS systems. Consider also synchronizing the startup of all IMS systems at a disaster recovery site.

If a sysplex failure or a restart at a disaster recovery site occurs, any new authorizations to databases or DEDB areas held by a failed IMS system are not granted until the failed IMS system is restarted. If each IMS system that failed is restarted one at a time and begins normal processing, authorization to databases or DEDB areas that were authorized by failed IMS systems will also fail. For full-function databases, message DFS047 is issued with a return code of X'09'. For DEDB areas, message DFS3709 is issued with a return code of X'09'.

After all of the failed IMS systems have been restarted, the databases and DEDB areas have to be manually started. Using the OPTION keyword and SYNCplex parameter on the /ERE command, you can serialize the restart of failed IMS systems. Issue the command on each IMS failed system.

After IMS restart processing completes, write-to-operator-with-reply (WTOR) message DFS3067A is issued. When all of the IMS systems receive this message, the operator can reply to each WTOR. Normal processing resumes, and database and DEDB area authorizations can be granted.

The OPTION SYSPLEX keyword and parameter are valid only on a /ERE command.

## **Authorizing databases after an IMS failure**

In a data-sharing environment, if an IMS system fails, another IMS system attempting to authorize a database or DEDB area held by the failed IMS system will fail.

For full-function databases, message DFS047 is issued with a return code of X'09'. For Fast Path DEDBs, message DFS3709 is issued with a return code of X'09'. The X'09' return code indicates that the database or area was previously authorized in one or more online IMS systems. The IRLM exit verification indicates that the IMS systems are inactive. Therefore, the identified database or DEDB area cannot be authorized.

To prevent this from occurring, specify the IRLM parameter SCOPE=NODISCON. This indicates that the IRLM is in a data-sharing environment, and that intersystem sharing is to be performed. The IRLM remains connected to the data-sharing group, even when no IMS systems are identified to it.

If an IMS system fails, the IRLM remembers the status of the failed IMS. Any new authorizations to databases that are held by the failed IMS will be allowed. The IRLM exit verification processing that occurs at restart recognizes the status of the failed IMS system and grants new authorizations.

---

## Chapter 16. Planning for VTAM generic resource groups

This topic describes the planning and administration tasks associated with using VTAM generic resource groups.

---

### Requirements for using VTAM generic resource groups

There are several requirements for balancing sessions using VTAM generic resource groups.

These requirements include:

- A sysplex environment.
- A z/OS coupling facility structure named ISTGENERIC defined in the active CFRM policy for the sysplex environment.
- A connection between the VTAM APPN node and the ISTGENERIC coupling facility structure.
- All IMS subsystems participating in the generic resource group reside in the sysplex environment.
- Each IMS subsystem participating in the generic resource group shares the same IMS generic resource name.
- An APPC generic resource name is defined to z/OS for LU 6.2 communications. Define this name on the GRNAME parameter of the LUADD statement in the APPC/MVS APPCPMxx member.
- All IMS systems that belong to the generic resource group must be defined with equivalent specifications.

**Related reading:** For information on VTAM and defining the ISTGENERIC coupling facility structure, see *z/OS Communications Server: SNA Network Implementation Guide*.

---

### VTAM generic resource group restrictions

When creating a VTAM generic resource group, system programmers should be aware of the restrictions.

Restrictions include:

- A particular IMS cannot be a member of more than one VTAM generic resource group.
- The generic resource name specified on the GRSNAME parameter must not match the name of an XRF USERVAR.

For more information on XRF and VTAM generic resources, see “XRF and VTAM generic resources” on page 292.

---

### VTAM generic resource affinity

After VTAM selects a specific IMS subsystem in a generic resource group to perform the work of a specific terminal, VTAM generic resource (VGR) affinity is established, and all of the terminal's subsequent sessions connect to the same IMS subsystem until IMS or VTAM resets the terminal's affinity.



IMS assigns VGR affinity management to IMS or VTAM for each session depending on the status recovery mode of that session. IMS ignores the GRAFFIN= parameter, which is obsolete for VGR. Multiple Systems Coupling (MSC) links use local affinity only.

**Note:** For ISC terminals and MSC links, VTAM and IMS do not release VGR affinities until all parallel sessions have been terminated.

## **VGR affinity management**

Either IMS or VTAM can manage VGR affinities. An important difference between the two is the ability of each to reset affinities for terminals in the event of an IMS failure. If a terminal or MSC link has an affinity with an IMS that fails, the terminal cannot continue work until its affinity is reset.

- IMS-managed – IMS tells VTAM when to delete VGR affinity when a session (or IMS) terminates. IMS deletes the affinity unless significant status for a terminal exists on a local IMS system.
- VTAM-managed – VTAM deletes VGR affinity when a session (or IMS) terminates.

When VTAM manages affinities, if IMS fails, VTAM can reset the affinities of terminals or MSC links that had sessions open with the failed IMS. With their affinities reset, the terminals can open a new session with another IMS in the generic resource group.

When IMS manages affinities, if IMS fails, IMS might not be able to reset the affinities between the terminals and the failed IMS. In this case, the terminals cannot open new sessions until the failed IMS is restarted. In the case of MSC links, the sessions cannot be moved to cloned links on another IMS when the affinity is set.

The type of terminal and the status recovery mode you set for the terminal determine whether IMS or VTAM manages VGR affinities.

## **IMS-managed affinity**

IMS manages VGR affinity for the following:

- Static terminals with recovery mode of LOCAL, where IMS sets VGR affinity because the terminal must not be allowed to access other IMS systems while significant status exists on the local IMS system. IMS ensures that VGR affinity matches RM affinity.
- Dynamic STSN terminals with status recovery mode of LOCAL, where IMS sets VGR affinity because the terminal must not be allowed to access other IMS systems while significant status exists on the local IMS system. IMS ensures that VGR affinity matches RM affinity because the user name and status recovery mode must be provided at logon.
- MSC links use local mode and IMS managed affinity only, and do not have RM support.

---

## Creating a VTAM generic resource group

You create a VTAM generic resource group by specifying the same generic resource name for all the IMS systems participating in the VTAM generic resource group.

You can specify the VTAM generic resource name for each IMS in two ways:

- The GRSNAME= startup parameter in the IMS and DCC procedures.
- The GRSNAME parameter of the /START VGR command when starting the control region of an IMS.

**Restriction:** The generic resource name specified on the GRSNAME parameter must not match the name of an XRF USERVAR.

### Related reading:

- For more information on the GRSNAME= startup parameter, see *IMS Version 11 System Definition*.
- For more information on the /START VGR command, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

---

## Specifying APPC generic resource names

Specifying an APPC generic resource name enables LU 6.2 devices to participate in the session balancing provided by a generic resource group. Specify the APPC generic resource name on either the GRNAME parameter of the LUADD statement in the APPC/MVS APPCPMxx member or with the SET APPC z/OS operator command.

**Restriction:** The APPC generic resource name must be different from the name you specify on the GRSNAME execution parameter.

---

## Initiating sessions with IMS in a VTAM generic resource group

Terminal sessions can be initiated either from IMS or from a terminal. All sessions initiated by IMS are generic sessions; that is, sessions in which affinity is established. MSC sessions are initiated with an /RSTART LINK command.

Sessions from a terminal can be initiated in the following ways:

- You can initiate a session with any IMS in a generic resource group by specifying the GRSNAME execution parameter. An affinity with the selected IMS is then established.
- You can initiate a session with a specific IMS by specifying the APPLID name of the system. No affinity is established.
- If an IMS in a generic resource group is also part of an XRF with MNPS system, you can initiate a session with that specific IMS using the MNPS name instead of the APPLID name. No affinity is established.

### Related reading:

- For more information on XRF and VTAM generic resources, see “XRF and VTAM generic resources” on page 292.
- For more information on XRF systems and specifying the MNPS name, see “Tailoring the IMS execution JCL” on page 646.

---

## XRF and VTAM generic resources

Extended Recovery Facility (XRF) systems help ensure session persistence and the availability of IMS resources. There are two types of XRF systems: XRF systems that use VTAM multinode persistent sessions (MNPS) and XRF systems that use a USERVAR. XRF with MNPS systems can participate in a VTAM generic resource group, but XRF with USERVAR systems cannot.

### XRF with MNPS and VTAM generic resources

For XRF with MNPS systems in a generic resource group, VTAM associates the generic resource name with the MNPS name of the XRF with MNPS systems instead of the APPLID name. Terminals that want to log on directly to an IMS that is part of an XRF with MNPS system must use the MNPS name.

Including an XRF with MNPS system in a generic resource group does not guarantee session persistence for terminals logging on using the generic resource name unless every IMS in the generic resource group is also part of an XRF with MNPS system.

If you use XRF with MNPS and are migrating off of the 3745 controller, be aware that you might encounter some performance limitations.

**Recommendations:** To minimize these performance limitations:

- Use the Communication Controller for Linux on System z® (CCL) as a replacement for the 3745 controller, which allows continued use of XRF.
- Migrate from XRF to VTAM Generic Resources (VGR), which requires a parallel sysplex environment.

### XRF with USERVAR and VTAM generic resources

Although XRF systems with USERVAR cannot participate in a generic resource group, they can coexist in the same sysplex as a generic resource group. In this case, generic resource names and USERVARs must be different.

XRF with USERVAR systems can participate in a shared-queues environment.

---

## Changing the logon procedure

You can provide the option of logging on using either a generic resource name or an IMS APPLID name; therefore, plan to make one or both of these names available in the logon procedure.

---

## Overriding the APPLID field

Each IMS identifies itself to VTAM using the application name you specify in the APPLID= keyword of the IMS COMM system definition macro and in the VTAM APPL definition statement.

In order to keep the IMS definitions as similar as possible, you can specify the same name in the APPLID= keyword of the COMM macros for all the IMS systems in a generic resource group. You must then override this name by specifying a unique APPLID name on the APPLID1 execution parameter.

Similarly, you can override the PASSWD field in the COMM macro by using the PASSWD= execution parameter.

**Related reading:** For more information on the APPLID1 and PASSWD execution parameters, see *IMS Version 11 System Definition*.

---

## Determining when affinities are terminated

When you terminate a session by logging off, IMS terminates the affinity between your node and the IMS VTAM generic resource member. However, in some situations, affinities persist.

You can use several methods in order to terminate these affinities.

Affinity for MSC and ISC parallel sessions is maintained as long as one or more sessions are active, or at least one session has significant status and the recovery mode is local. MSC sessions terminate the VTAM affinity when a /PSTOP LINK command is issued for all parallel link sessions. MSC recovery mode is always local.

ISC sessions terminate the VTAM affinity when:

- A /QUIESCE NODE command is issued (cold termination) for all sessions and recovery mode is local or global.
- A /CLOSE NODE command is issued (warm termination) for all parallel sessions and the recovery mode is global.

The VTAM affinity is also terminated when IMS is shut down normally by a /CHE FREEZE|PURGE|DUMPQ command and all sessions are successfully terminated. The /CHE FREEZE|PURGE|DUMPQ LEAVEGR command terminates all affinity unconditionally. Whether the shutdown is initiated by IMS (/CHE) or by the session (/PSTOP, /CLOSE, or /QUIESCE), the affinity is maintained if any parallel MSC or ISC sessions between the same IMS systems remain active or have been terminated but still have a significant status.

If there are parallel sessions using both MSC and ISC between the same IMS systems, then the previously mentioned affinity rules are applied simultaneously. That is, the affinity is maintained for both MSC and ISC parallel sessions if any link or node has significant status.

In order for IMS to terminate the affinity in all cases, the VTAM ACB must be opened with a /START DC command.

MSC significant status (MSC recovery mode is always local)

- Link is in active status
- Link is in ERE status

ISC significant status

- Node is in active status
- Node is using RM local recovery mode (SRMDEF=Local for the node) and is in warm terminate status

To display the list of IMS systems and their affinities with specific nodes or MSC links, use the /DISPLAY AFFIN command.

**Related reading:** For information on the /DISPLAY AFFIN NODE or /DISPLAY AFFIN LINK command, see *IMS Version 11 Commands, Volume 1: IMS Commands A-M*.

For information on how to terminate affinities that persist, see “Terminating affinities that persist.”

---

## Terminating affinities that persist

A terminal cannot participate in session balancing if it has a persisting affinity with a particular IMS from a previous session. When VTAM manages the VGR affinity, affinity is automatically reset at session termination.

When IMS manages the VGR affinity, an affinity persists after a logoff or IMS failure in each of the following situations:

- The terminal is static with LOCAL status recovery mode, and has end-user significant status (Conversation, STSN, Fast Path, or full-function response mode).
- The terminal is ETO SLUP or 3600 Finance terminal with LOCAL status recovery mode and RCVYSTSN=YES.
- The terminal is an ISC parallel-session terminal with LOCAL status recovery mode that has not quiesced.
- The terminal is an ISC parallel-session terminal, and one or more of the parallel sessions has not quiesced.
- The MSC link is in an ERE status.

You can terminate an affinity using any of the following methods:

- Enter the /CHECKPOINT command with the LEAVEGR keyword.
- Use the Logoff or Signoff exit routine to reset terminal status.
- Cold start the IMS TM subsystem, and specify the GRSNAME when restarting IMS.

### **Related reading:**

- For information on the /CHECKPOINT command and the LEAVEGR keyword, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.
- For information on the Logoff and Signoff exit routines, see *IMS Version 11 Exit Routines*.

---

## Dropping an IMS system from a generic resource group

You can use the /STOP VGRS command to drop an IMS out of a generic resource group. Similarly, you can use the /START VGRS command to bring that IMS back into the generic resource group.

---

## Resetting terminal status

You can reset terminal status (and thus terminate affinities) using commands, exit routines or parameters.

The following list includes the ways that you can reset the terminal status:

- An authorized command
- The Logoff (DFSLGFX0) exit routine

- The Signoff (DFSSGFX0) exit routine
- The RCVYFP, RCVYRESP, RCVYSTSN, and RCVYCONV parameters
- The NONE status recovery mode

The Logoff (DFSLGFX0) exit routine resets a terminal's status during a logoff. Similarly, you can use the Signoff (DFSSGFX0) exit routine to reset a user's status during a signoff. Resetting terminal or user status enables IMS to terminate an affinity with the terminal, enabling the terminal to log on again and participate in session balancing.

#### **Related reading:**

- For more information on the Logoff and Signoff exit routines, see *IMS Version 11 Exit Routines*.
- For more information on the NONE status recovery mode, see “Status recovery mode for end-user significant status” in *IMS Version 11 Communications and Connections*.
- For more information on the RCVYFP, RCVYRESP, RCVYSTSN, and RCVYCONV parameters, see *IMS Version 11 System Definition*.

---

## **Logging on after IMS failure with affinity**

Suppose you are operating in a shared-queues environment, and your terminal is logged onto an IMS in a generic resource group. If that IMS fails, the affinity between that IMS and your terminal might persist.

Consequently, you can have output messages waiting for you on that terminal. To obtain your messages, you can log onto another IMS within the generic resource group, provided that all of the following conditions exist:

- You log on using the APPLID name of another IMS within the generic resource group.
- Your output messages are not locked on the failed IMS, and those messages are not responses to either an IMS conversation or a Fast Path response-mode transaction.
- If you have a resource structure and Resource Manager, the status recovery mode is not LOCAL and there is not end-user significant status on the failed IMS (no RM affinity).

---

## **Controlling VTAM generic resource member selection**

You can control VTAM's selection of a generic resource member for a terminal when the terminal logs on. To control the selection process, use either the VTAM Generic Resource Resolution exit routine (ISTEXCGR) or the z/OS Workload Manager.

Choosing between these facilities depends on the needs of your application program and on your installation requirements.

The criteria that VTAM uses to select a generic resource member are:

1. Existing affinity. VTAM maintains an affinity table in a list structure called ISTGENERIC on the coupling facility.
2. VTAM Generic Resource Resolution exit routine (ISTEXCGR).
3. z/OS Workload Management.
4. Current session counts.

### **Related reading:**

- For more information on the VTAM Generic Resource Resolution exit routine (ISTEXCGR), see *z/OS Communications Server: SNA Customization*.
- For more information on the z/OS Workload Management, see *z/OS MVS Planning: Workload Management* and *z/OS MVS Programming: Workload Management Services*.

---

## **Ensuring consistency across the IMSplex**

Inconsistencies in definitions, operator commands and procedures, and exit routines can create problems within a generic resource group.

- At a minimum, they can confuse terminal users, who are unaware of being in session with different IMS systems from one logon to another logon.
- At worst, they can cause processing disruptions.

**Restriction:** LU6.1 parallel sessions between IMS and another node must all use VGR or non-VGR. LU6.1 parallel VGR and non-VGR sessions (ISC or MSC) cannot be mixed because the first parallel session established between the LU6.1 partners establishes how the partner node is known, by either the real APPLID name (non-VGR) or the GRSNAME (VGR). This condition is remembered by VTAM, and VTAM sets a flag (INRIFLG0 = NIBNNAMS in the ISTNRIPL parameter area) on subsequent parallel session initiations to indicate this name association. IMS uses this flag to establish the session. The mismatch could cause the session initiation to fail (such as DFS3645 Logon Rejected) or the session affinity to be incorrectly managed by IMS.

You can mix VGR and non-VGR between non-parallel LU6.1 sessions, such as simultaneous use of IMS1 to IMS2 using VGR, IMS1 to CICS using non-VGR, or IMS1 to IMS3 using non-VGR.

**Recommendations:** To ensure consistency across a generic resource group:

- Use equivalent specifications when defining the member IMS systems.  
**Example:** If one IMS specifies that ETO be included (ETOFEAT=YES on the IMSCTRL macro), all other IMS systems in the generic resource group should specify that ETO be included.
- Specify execution parameters consistently across generic resource members.  
**Example:** IMS systems in a generic resource group should each specify that ETO be enabled or not be enabled.
- If ETO is included, specify the same information on the ETO descriptors for each IMS generic resource group member.
- Use the same naming conventions for LTERMs and transactions across a generic resource group. If the same LTERM name is assigned to different nodes on different IMS systems, unpredictable results might occur.

**Related Reading:** For more information about how to ensure name uniqueness and resource type consistency, see Chapter 17, “Planning for Transaction Manager resources in an IMSplex,” on page 299.

- Make all terminal definitions consistent across generic resource members.  
**Example:** LTERM names should be consistently defined, during IMS system definition, to the same physical terminals on all the IMS systems in the generic resource group.
- MSC links in a generic resource group should be cloned.



**Example:** Generic resource groups should have the same MSPLINK, MSLINK, and MSNAME characteristics

- Ensure that master terminal operators (MTOs) on each IMS in the generic resource group have procedures for notifying the other MTOs about the commands they issue. If you are using a single point of control (SPOC), commands go to all IMS systems that can process the commands and there is no need for the notification procedures.

**Example:** If you do not have Resource Manager (RM) and one MTO issues a /STOP command to stop a user on one IMS in a generic resource group, that user is stopped only on the IMS on which /STOP is entered. If you have RM and a resource structure, however, RM manages the user globally and the user is stopped on all the IMS systems in the IMSplex.

**Related reading:** For more information on using a SPOC, see “Controlling IMS with the TSO SPOC application” in *IMS Version 11 Operations and Automation*.

- Make exit routines functionally equivalent across a generic resource group—especially the Destination Creation exit routine, the OTMA Routing exit routines, the Logon and Signon exit routines, and the Logoff and Signoff exit routines.

---

## Bypassing affinity management during the IMS ESTAE process

If you use IMS to manage generic resource affinities, you can control whether or not IMS uses or bypasses the VTAM generic resource logic in the IMS ESTAE exit. To use the IMS ESTAE process, indicate GRESTAE=Y for terminal sessions and GRMESTAE=Y for MSC link sessions in the DFSDCxxx IMS.PROCLIB member data set.

GRESTAE=Y or GRMESTAE=Y indicates that IMS should follow the existing ESTAE logic to delete affinity for nodes where no status remains before closing the ACF/VTAM ACB. GRESTAE=N or GRMESTAE=N indicates that IMS should close the ACF/VTAM ACB immediately to expedite IMS termination, and leave affinity for all nodes set.

IMS VTAM generic resource logic in the IMS ESTAE exit attempts to delete generic resource affinity if no terminal or MSC link status remains. The IMS ESTAE routine uses the VTAM delete affinity force option (ENDAFFNF) which does not require the sessions to be disconnected from the VTAM terminal first.

The following table summarizes how generic resource affinities are managed based on whether affinities are VTAM or IMS managed and which GRESTAE options are indicated in the DFSDCxxx IMS.PROCLIB member data set.

*Table 21. Generic resource affinities management and GRESTAE options*

| Affinity manager | GRESTAE/<br>GRMESTAE<br>option | Generic resource affinities management                                       |
|------------------|--------------------------------|------------------------------------------------------------------------------|
| IMS              | Y                              | IMS manages generic resource affinities, including during ESTAE processing.  |
|                  |                                | IMS resets affinity during the ESTAE process that occurs during IMS failure. |



*Table 21. Generic resource affinities management and GRESTAE options (continued)*

| Affinity manager | GRESTAE/<br>GRMESTAE<br>option | Generic resource affinities management                                               |
|------------------|--------------------------------|--------------------------------------------------------------------------------------|
| IMS              | N                              | IMS manages generic resource affinities, except during ESTAE processing.             |
|                  |                                | IMS does not reset affinity during the ESTAE process that occurs during IMS failure. |
| VTAM             | Y or N                         | VTAM manages generic resource affinities.                                            |

---

## Chapter 17. Planning for Transaction Manager resources in an IMSplex

This topic describes in detail the functions of the resource structure and Resource Manager (RM) as they relate to managing IMS Transaction Manager resources in an IMSplex.

The resource structure and RM are intended to run in a DB/DC or DCCTL environment, and with CSL operating.

### **Related reading:**

- For more information on administering an IMSplex with CSL, see Chapter 7, “CSL administration,” on page 141.
- For information on resource-structure-related parameters for the DFSDCxxx member of the PROCLIB data set, see *IMS Version 11 System Definition*.
- For information on user exit routines DFSINSX0, DFSSGNX0, DFSINTX0, and DFSLGNX0, see *IMS Version 11 Exit Routines*.
- For a list of the commands that relate to the resource structure and RM, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

---

## Resource name uniqueness

IMS ensures that a resource name is active only once in the IMSplex at any particular time. IMS systems within the IMSplex cannot activate the same resource at the same time. The IMSplex automatically enforces resource name uniqueness only when Resource Manager (RM) is active and a resource structure is defined in the coupling facility.

When a resource is active, it is owned by one IMS in the IMSplex. No other IMS can activate the same resource until it becomes inactive on the owning IMS.

Name uniqueness is enforced only for VTAM LTERMs, VTAM single-session nodes, user IDs, and users. IMS no longer supports BTAM terminals.

Name uniqueness is enforced for user IDs only if SGN is not M. User ID name uniqueness enforcement is not automatic. However, name uniqueness for VTAM LTERMs, VTAM single-session nodes, and users is automatically enforced if you have RM and a resource structure defined.

For ISC terminals, if the option for disabling resource sharing in the IMSplex is on in the DFSINTX0 user exit routine, resource name uniqueness is not enforced for static LU 6.1 sessions.

### **Disabling enforcement of resource name uniqueness**

You can disable the enforcement of resource name uniqueness, and TM resource sharing in general, by specifying STM=NO in the DFSDCxxx PROCLIB member data set.

**Related reading:** See “Transaction Manager resources” for more detailed information about how and why IMS enforces name uniqueness for particular resources.

---

## Resource type consistency

Resource type consistency ensures that a name is unique within a group of resources, called a *name type*. The IMSplex automatically enforces resource type consistency when RM is active and a resource structure is defined in the coupling facility.

IMS enforces resource type consistency for the name type of message destination. IMS ensures that a resource defined as a message destination is consistent across the IMSplex.

Message destinations are the following:

- LTERMs
- LTERMs defined as APPC descriptors
- MSNAMEs
- IMS-defined transactions
- CPI-C transactions

### Disabling enforcement of resource type consistency

You can disable the enforcement of resource type consistency when a resource structure is present by specifying STM=NO in the DFSDCxxx PROCLIB member data set. After specifying STM=NO, resource type consistency is enforced only for IMS-defined and CPI-C transactions.

**Related reading:** See “Transaction Manager resources” for more detailed information about how and why IMS enforces type consistency for particular resources.

---

## Global callable services

Callable services are provided for user-provided exit routines in order to find resources such as nodes, LTERMs, and users. Callable services returns global resource information shared in the resource structure. If no global information is available, local information is returned by default.

If you want only local information, use the input flag LOCAL on the CSCBLK DSECT, which is defined by the DFSCCBLK macro.

---

## Transaction Manager resources

This topic discusses the following types of Transaction Manager (TM) resources.

### TM resources: APPC descriptors

IMS defines APPC descriptors to RM at initialization or during /STA LU62DESC to maintain resource type consistency for message destinations.

IMS ensures that an APPC descriptor name is not used as a transaction, MSNAME, or LTERM name. Descriptors can be defined on multiple IMS systems as name uniqueness is not enforced.

## APPC descriptor creation and deletion

At initialization or during /STA LU62DESC, IMS ignores any descriptors already defined as a transaction, MSNAME, or LTERM and issues a warning message. APPC descriptors are deleted from RM when all IMS systems that have defined the descriptor have terminated, warm or cold. Systems that do not define the descriptor do not need to terminate.

## TM resources: VTAM LTERMs

IMS defines VTAM LTERMs to RM for resource type consistency for message destinations, name uniqueness for LTERMs, and LTERM status recovery.

### LTERM name uniqueness

IMS enforces name uniqueness for VTAM LTERM names to ensure output security within an IMSplex. When an LTERM becomes active, IMS ensures that the LTERM does not become active on any other IMS and is unavailable to any other user or terminal.

For LU 6.1 sessions, LTERM name uniqueness can be ignored by using the DFSINTX0 user exit routine. See *IMS Version 11 Exit Routines* for more information about this exit routine.

IMS no longer supports BTAM terminals.

### LTERM creation and deletion

When an LTERM becomes active during signon or logon, IMS defines it to the RM. If the name is the same as a transaction, APPC descriptor, or MSNAME, or if the LTERM is active on another IMS, the LTERM does not become active. The signon of a dynamic user continues without an LTERM assigned (but if there is only one LTERM, the signon is rejected), and the logon for a static terminal is rejected. IMS also defines LTERMs to RM when recoverable commands affecting significant status are processed.

When an LTERM becomes inactive and there is no significant status data, IMS deletes the LTERM. An LTERM is deleted in three situations:

- During a logoff or signoff
- When a command deleting significant status is processed
- During resource cleanup after an IMS failure and no significant status exists

## TM resources: MSNAMEs

MSC networks use MSNAMEs to define remote IMS systems and logical link paths between remote and local IMS systems in an MSC network. For MSNAMEs, Resource Manager (RM) enforces only resource type consistency for message destinations, but not resource name uniqueness.

When IMS initializes, IMS defines to RM each MSNAME that has a remote SYSID. Because resource name uniqueness is not enforced, the same MSNAMEs can be defined on multiple IMS systems.

## Creation and deletion of the MSNAME TM resource

An MSNAME becomes a TM resource to RM at IMS initialization. RM saves the MSNAME in the resource structure. MSNAMEs are not deleted from a resource structure once they are defined and cannot be reused unless the resource structure itself is deleted.

**Related reading:** For additional information on MSNAMEs and how they are used in a shared resources group, see *IMS Version 11 Communications and Connections*.

## TM resources: VTAM terminal nodes

IMS defines nodes to RM to enforce name uniqueness for single-session VTAM terminals and to recover node status.

### Node name uniqueness

IMS enforces name uniqueness of VTAM terminals to ensure only one terminal can be logged on at one time and to ensure data integrity. These terminals include all VTAM terminals except LU 6.1 (ISC) parallel sessions.

IMS does not enforce name uniqueness for LU 6.1 parallel sessions, but data integrity is ensured by user name uniqueness. For LU 6.1 sessions, node name uniqueness can be ignored by using the DFSINTX0 user exit routine. See *IMS Version 11 Exit Routines* for more information on this exit routine.

**Note:** Some session manager products that assign non-unique ETO node names to terminals logging onto multiple systems in an IMSplex will not work properly with RM. IMS prevents the same non-ISC node name from being active in more than one system at a time.

### Node creation and deletion

IMS defines a node to RM when the node becomes active during terminal logon. Logon is rejected if the node is already active on another IMS system, except for ISC parallel sessions. IMS also defines a node to RM when recoverable commands that process significant status are processed. These commands include the commands that set global MFSTEST, STOP, and TRACE.

When a node becomes inactive and there is no significant status data, IMS deletes the node. A node is deleted in three situations:

- During logoff
- When a command deleting significant status is processed
- During resource cleanup after an IMS failure and no significant status exists

## TM resources: transactions

IMS defines a transaction dynamically with the type-2 CREATE TRAN command or statically in the system where the application will run or dynamically as a CPI-C transaction executed by an APPC conversation. A transaction can run in multiple systems concurrently as name uniqueness is not enforced.

Within the context of TM resource management, transactions are treated as a resource only to prevent another TM resource from duplicating the transaction

name. RM continues to prevent the duplication of transaction names even if you disable TM resource sharing by specifying STM=NO in the DFSDCxxx PROCLIB member data set.

## Transaction resource type consistency

IMS ensures that a transaction name cannot be used as an MSNAME, APPC descriptor, or an LTERM name. IMS ignores a CPI-C transaction if the name is defined as a different message destination. APPC still validates the destination and does not allow LTERMs as destinations.

## Transaction creation and deletion

IMS defines a transaction to RM when IMS initializes, when online change adds a transaction, or when a transaction is created with the type-2 CREATE TRAN command. IMS ignores the transaction if the name is already defined in the IMSplex as another resource.

IMS identifies a CPI-C transaction to RM during APPC input processing and after it checks to make sure the name is not defined as another message destination already. If it is defined as another resource, IMS uses the defined destination and no error is generated.

Cold-starting an IMS does not delete a transaction.

## TM resources: user names

The user name and the user ID are usually the same; however, user exits and descriptors can override the user name. The *user* is the user signed on to a dynamic terminal or parallel session subpool and has associated work and status. The *user ID* identifies a person signed on to a terminal for security authorization by a security product such as RACF.

## User name uniqueness

IMS enforces name uniqueness for user names to ensure data integrity for users. Users are defined to RM during user signon or ISC session initiation. If the user is already active, IMS rejects the signon or session initiation.

For ISC sessions, user name uniqueness can be ignored by using the DFSINTX0 user exit routine. See *IMS Version 11 Exit Routines* for more information about this exit routine.

## User creation and deletion

When a user becomes active during user signon, IMS defines it to the RM. If the user is active on another system, signon is rejected. IMS also defines users to RM when recoverable commands affecting significant status are processed.

IMS deletes the user from RM when the user becomes inactive and no significant data exists. A user is deleted in three situations:

- During user signoff
- When a command deleting significant status is processed
- During resource cleanup after an IMS failure and no significant status exists

## TM resources: user IDs

The *user ID* identifies a person signed on to a terminal for security authorization by a security product such as RACF.

The user ID and the user name are usually the same, but see “TM resources: user names” on page 303 for an example of when they are not the same.

### User ID name uniqueness

You can decide whether to enforce name uniqueness for user IDs. The default is to enforce user ID name uniqueness. If you want one user ID to be able to sign on to multiple terminals at a time (name uniqueness not enforced), use the IMS startup parameter SGN=M.

### User ID creation and deletion

When a user signs on to a terminal and single-signon enforcement is requested, IMS defines the user ID to RM. Once a user is signed onto a VTAM terminal in the IMSplex, any signon attempt by that user is rejected. When a user ID becomes inactive, IMS deletes it from RM. A user ID is deleted in two situations:

- During user signoff
- During resource cleanup after an IMS failure

---

## How RM and the resource structure impact IMS activities

This topic provides information on how RM and the resource structure impact some IMS activities.

### Conversations

Conversation status is shared within the IMSplex, and so the conversation's IDs must be unique for individual dynamic users or static nodes. To uniquely identify a conversation, qualify the conversation with the associated input LTERM.

To display conversation status, use the `/DISPLAY CONVERSATION` command for local and RM conversation information.

### Initialization

IMS enforces consistency of single signon, SGN=M, in the startup parameters when RM and a resource structure are active. The first IMS active in the IMSplex sets the value for the entire IMSplex. If an IMS joins an IMSplex with a different value for single signon, IMS issues a warning message and uses the IMSplex value. You can change the specification when all the IMS systems leave and then rejoin the IMSplex (warm start). You can override the value by using the `/NRE` or `/ERE` commands.

During a warm start, when RM is active, the global signon value comes from the startup parameter of the first IMS to join the IMSplex. The signon value comes from the checkpoint log records when RM is inactive during the warm start.

## Shutdown

When an IMS system shuts down, all nonrecoverable status owned by the system is deleted from the resource structure. SCI notifies the other IMS systems that the IMS is leaving because of either a normal or abnormal shutdown.

When an IMS system shuts down abnormally, another IMS in the IMSplex becomes the cleanup IMS. This IMS requests a list of resources without significant status. If a resource does have significant status and the resource recovery mode is GLOBAL, the owner is cleared and the resource becomes available to other systems. If the recovery mode is LOCAL, affinity for the failed system is enforced.

## Checkpoint

Because terminal and user status are maintained by RM, if active, local DC control blocks must be cleaned up and deleted during IMS checkpoint if they are inactive. Checkpoint deletes local status in local blocks with GLOBAL status recovery mode. Checkpoint maintains status in local blocks with LOCAL status recovery mode.

These recovery modes are treated the same way during a simple checkpoint after a warm or emergency restart. However, if a resource is in LOCAL status recovery mode, DFSSGNX0 and DFSLGNX0 can request that another IMS steal the resource that is owned by the failed IMS system. In that case, the restarting IMS deletes the local status during its restart.





---

## Chapter 18. IMS security

This topic provides information to help you establish resource security for an IMS online system and identifies the resources that can be protected and the facilities available to protect them, provides design considerations, and describes the steps that you need to take to activate security.

Although there are many similarities and overlaps, the security for the DB/DC and DCCTL environments is discussed separately from the security for the DBCTL environment.

### **Related concepts:**

“ODBM and security” on page 150

---

### Overview of DB/DC and DCCTL security

When you initiate security safeguards, you must balance requirements between those responsible for the security of resources and those users who legitimately need access to those resources.

Because the person who is assigned to resource security is held responsible for resources that might be compromised, that person should not allow easy access to dominate protection measures. However, users performing their assigned tasks need convenient access to the resources. The users and the security specialist should work out a balanced approach between the ease of resource access and the complexity of protecting that resource.

In an IMS system, two types of security exist. You can address one or both types:

- Securing the kind of resource to which a user has access. For example, a user might be allowed access to the Part database but not to the Customer Order database.
- Securing what the user can do to the resource after that user has access to it. For example, a user might be allowed to read a file but not to update it.

### **DB/DC and DCCTL resources that can be protected**

Before you decide what security facilities to use in designing a secure IMS system, you should know which resources within the system need protection.

In other words, you should decide what to protect before you decide how to protect it.

The following resources can be protected in the DB/DC and DCCTL environments:

- IMS control region
- IMS online system
- IMS system data sets
- Transactions
- Commands
- Program specification blocks (PSBs)
- Online application programs
- Databases

- Dependent regions
  - BMP (batch message processing) region
  - IFP (IMS Fast Path) region
  - JBP (Java batch processing) region
  - JMP (Java message processing) region
  - MPP (message processing program) region
- Terminals
  - Logical terminals (LTERMs)
  - Physical terminals (PTERMs)
  - Master terminals

## Defining security during DB/DC and DCCTL system definition

You can make IMS security choices with initialization EXEC parameters and in three system definition macros: SECURITY, COMM, and IMSGEN. These specifications let you choose the type of security that is active during online execution.

You can make other security choices, including which resources you want to protect, by using the Resource Access Control Facility (RACF). You can also specify security choices using the security parameters in the IMS and DCC startup procedures.

If you do not specify any security in any of the three system definition macros, IMS provides a basic level of resource security called *default security*, which:

- Prohibits the entry of certain commands from any terminal other than the master terminal. This basic security function is activated upon completion of stage 2 of IMS system definition. When you implement input-access security with RACF, IMS removes the default security restrictions. In the case of static terminals if sign-on is not required, IMS uses the control region user ID for command validation.
- Applies only to statically defined terminals. Terminals that are defined by using ETO are automatically governed by an identical level of default security. When you modify and use the Command Authorization exit routine (DFSCCMD0), IMS removes the default security for dynamically defined terminals.

All security specifications are consolidated in the SECURITY macro. However, you can also specify security options in the COMM macro, IMSGEN macro, or with initialization EXEC parameters. The security-related specifications are accepted hierarchically in the following order:

1. SECURITY or initialization EXEC parameters
2. COMM
3. IMSGEN

Security specifications in the SECURITY macro or initialization EXEC parameters override security specifications that are coded in the COMM and IMSGEN macros.

## Security facilities for DB/DC and DCCTL resources

The following table summarizes the resources you can protect, the types of security that you can use to protect the resources, and the security facilities that you can use to implement each type of security.

When choosing which security facilities to use, also consider your installation's security standards and operating procedures.

*Table 22. DB/DC and DCCTL resources and the facilities to protect them.*

| Resource                                              | Type of security                                                  | Security facility                          |
|-------------------------------------------------------|-------------------------------------------------------------------|--------------------------------------------|
| IMS control region and online system                  | Extended resource protection (using APPL resource class)          | RACF                                       |
| System data set                                       | z/OS password protection                                          | z/OS                                       |
|                                                       | Data set protection (VSAM)                                        | RACF                                       |
| Database                                              | Segment sensitivity                                               | PSBGEN RACF                                |
|                                                       | Field sensitivity                                                 | PSBGEN RACF                                |
|                                                       | Password security (for the /LOCK and /UNLOCK commands)            | RACF                                       |
| PTERM                                                 | Signon verification security                                      | RACF with exit routine                     |
|                                                       | Terminal-user security                                            | RACF                                       |
|                                                       | Password security (for the /IAM, /LOCK, and /UNLOCK commands)     | RACF                                       |
| LTERM                                                 | Password security (for the /IAM, /LOCK, and /UNLOCK commands)     | RACF                                       |
|                                                       | Application group name security                                   | RACF                                       |
|                                                       | Resource Access Security                                          | RACF                                       |
| Terminal defined with ETO                             | Signon verification security                                      | RACF and exit routine                      |
|                                                       | Input access security                                             | RACF and exit routine                      |
| LU 6.2 inbound and IMS-managed outbound conversations | Allocate verification security                                    | RACF and exit routine                      |
|                                                       | Input access security                                             | RACF and exit routine                      |
| PSB                                                   | Application group name security                                   | RACF                                       |
|                                                       | Resource Access Security                                          | RACF                                       |
|                                                       | APSB SAF security                                                 | RACF <sup>1</sup>                          |
| Transaction                                           | Application group name security                                   | RACF                                       |
|                                                       | Input access security                                             | RACF                                       |
|                                                       | Resource Access Security                                          | RACF                                       |
|                                                       | Password security (for the /LOCK and /UNLOCK commands)            | RACF                                       |
| Command                                               | default security                                                  | System definition                          |
|                                                       | Transaction command security for automated operator (AO) commands | RACF or Command Authorization exit routine |
|                                                       | Input access security                                             | RACF                                       |
|                                                       | type-2 command security                                           | RACF                                       |
|                                                       | DBRC command authorization <sup>2</sup>                           | RACF or exit routine                       |
| Type-1 Automated Operator Interface applications      | Transaction command security                                      | RACF or Command Authorization exit routine |
| Type-2 Automated Operator Interface applications      | Transaction command security                                      | RACF or Command Authorization exit routine |

Table 22. DB/DC and DCCTL resources and the facilities to protect them (continued).

| Resource                                                                                                                                                                                                           | Type of security                                              | Security facility |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------|-------------------|
| Online application program                                                                                                                                                                                         | Password security (for the /IAM, /LOCK, and /UNLOCK commands) | RACF              |
|                                                                                                                                                                                                                    | Extended resource protection (using APPL keyword)             | RACF              |
| Dependent region                                                                                                                                                                                                   | Application group name security                               | RACF              |
|                                                                                                                                                                                                                    | APSB SAF security                                             | RACF              |
|                                                                                                                                                                                                                    | Resource Access Security                                      | RACF              |
| <b>Notes:</b>                                                                                                                                                                                                      |                                                               |                   |
| 1. Using RACF to secure APSBs applies only to CPI-C driven applications and ODBA applications.                                                                                                                     |                                                               |                   |
| 2. DBRC command authorization is an additional command security option that applies only to DBRC commands. DBRC commands are also subject to any other command security options that are active in the IMS system. |                                                               |                   |

## Designing security for IMS DB/DC and DCCTL

This topic explains how the various types of IMS security can be used. When you are considering each part of your security design, consider the physical actions that a user must take to obtain access to the system. You probably will use more than one type of security checking.

This topic assumes the following control points for security:

- User identifications
- The master terminal
- Automated operator programs
- The use of regions

This section also assumes that you will use the Resource Access Control Facility (RACF) security product, which is licensed with the z/OS operating system. RACF is external to IMS.

### Limiting access from a terminal

To control access for a terminal control point, have the end users identify themselves each time they use a terminal and authorize the set of transactions and commands each user needs.

For a static terminal, or a dynamic terminal that has the SPQBname the same as the node name, a user will not be allowed to signon unless all conversations are held, or the user is authorized to use the transaction for the active conversation.

If there is an active conversation for a static terminal, and the user is not authorized to use its transaction, the user can enter a /HOLD command prior to signing on to put all of the conversations in a held state. The user will then be allowed to sign on.

If there is an active conversation for a dynamic terminal that has the SPQBname the same as the node name, only a user that is authorized to use the transaction of the active conversation will be allowed to sign on. The /HOLD command is not allowed prior to signing on for a dynamic terminal.

## Controlling access by signon verification

Signon verification security requires the entry of a user ID as a parameter on a /SIGN command at all terminals or a subset of the terminals.

To implement signon verification for statically defined VTAM terminals, set the OPTIONS= keyword to SIGNON in the TYPE or TERMINAL system definition macro. The specifications in the TERMINAL macro override those set in the TYPE macro. The default for OPTIONS= in the TYPE macro is NOSIGNON.

To avoid setting the OPTIONS= keyword in multiple system definition macros when you want all of your static terminals to sign on, specify SIGNON=ALL in the DFSDCxxx PROCLIB member. This requires all terminals to sign on except MTO, LU6.1, 3284/3286, and SLU1 printer-only devices.

Signon verification is done by RACF, an exit routine, or both. If you are using RACF, RACF checks for user ID, password, and group. If you are using an exit routine to implement signon verification, the exit routine can check for user ID and password.

To enable password verification:

- Set the password verification byte on the Initialization exit routine (DFSINTX0)
- Reply to the DFS3656 message
- Issue the /SIGN ON VERIFY command

The exit routines that you might use to verify user IDs and passwords include the Sign On/Off Security exit routine (DFSCSGN0). For ETO terminals, you might also use the Sign-On exit routine (DFSSGNX0). DFSCSGN0 is described in more detail in “Preparing security exit routines” on page 331.

**Recommendation:** If you have specified mixed-case password support in your IMS system definition (the PSWDC= parameter on the DBC, DCC, or IMS procedures), make sure that any exit routines that manipulate passwords can also support mixed-case passwords.

**Related reading:** For more information about these and other security exit routines, see *IMS Version 11 Exit Routines*.

## Using RACF PassTickets

A RACF PassTicket is a one-time-only password that is generated by a requesting product or function. The PassTicket is an alternative to the RACF password and removes the need to send RACF passwords across the network in clear text.

PassTickets also make it possible to move the authentication of a mainframe application user ID from RACF to one of these alternatives:

- Another authorized function executing on the host system
- The workstation local area network (LAN) environment

PassTickets are resistant to reuse because they only give one user access to a specific application for approximately 10 minutes. For most applications, once a particular PassTicket is used, the same user cannot use it again for the same application during the same 10-minute interval.

When a /SIGN ON command is received by IMS that contains a PassTicket instead of a password, the signon process fails, unless the PassTicket was created using the IMSID as the application name (name of the PTKTDATA profile). Because the

IMSID may not be known to other systems that might enter the signon command, flexibility is provided to allow for the use of other names as application names when creating PassTickets. The keyword APPL in the /SIGN ON command allows the end-user or program to specify another name (such as the IMS VTAM application name) rather than the IMSID when creating the PassTicket.

In a VTAM generic resource (VGR) environment, the remote end-user does not know which IMS will be chosen for the connection. The DFSDCxxx PROCLIB member provides a system-wide default application name (replacing IMSID) for all the IMS systems in a generic group. With the use of this replacement name for the application name, the creator of the PassTicket does not have to know which IMS system will be the recipient of the IMS terminal session.

#### **Related reading:**

- For more information about PassTickets, see *z/OS Security Server RACF Macros and Interfaces*.
- For more information on enabling the use of PassTickets, see *z/OS Security Server RACF Security Administrator's Guide*.

#### **Creating and using a RACF PassTicket:**

When you call RACF, a general PassTicket generator algorithm uses special input information to create a unique PassTicket. The PassTicket is an 8-character alphanumeric string that contains the characters A through Z and 0 through 9.

The algorithm generates a PassTicket using the following four input values:

- A RACF host user ID
- The PTKTDATA class profile name (also known as the application id)
- The RACF Secured Signon application key, which is contained in the PTKTDATA profile
- The required time and date input data from the application that is providing the logon function

**Requirement:** For RACF to properly evaluate PassTickets, the TOD clock must be set to GMT rather than local time.

## **Authorizing transactions and commands**

This topic describes transactions and commands, and explains how to authorize them.

### **Authorizing transactions**

Transaction authorization determines if a user ID is permitted to use a certain transaction. You can use an exit routine, RACF, or both to perform the transaction authorization. To gradually phase in RACF as your installation's transaction authorization method, use the Transaction Authorization exit routine (DFSCTRN0).

This routine can reject the transaction if the transaction is entered by an ETO terminal but is not protected by RACF. If you use both DFSCTRN0 and RACF, DFSCTRN0 is effective only after RACF has authorized the transaction or when the transaction is not defined to RACF. DFSCTRN0 is described under "Preparing security exit routines" on page 331.

If you specify the REVERIFY option to RACF, the user must reenter the signon password with each transaction code. REVERIFY is not supported when a takeover occurs in an XRF complex. Users might need to sign on again after a takeover if they are not on a class-1 terminal.

With program-to-program switching through the DL/I change (CHNG) call or by changing the transaction code in the SPA, you can use RACF, an exit routine, or both, to check transaction authorization. The same applies when the transaction code status is changed by the /SET, /LOCK, and /UNLOCK commands. In addition, as the application program associated with the transaction produces database changes, the user ID is logged with the change records on the IMS system log to identify the changes performed by a specific user.

IMS provides the RVFY= parameter in the IMS procedure for customers who want to force reverification that the operator who signed on to a terminal is the same operator who is now entering a command or transaction. This reverification is done with RACF by including the word 'REVERIFY' in the APPLDATA field of the command or transaction profile. For example:

```
RDEFINE Txxx tran-name UACC(NONE) APPLDATA('REVERIFY')
```

Each time the user enters this transaction code, the RACF password must be entered where an IMS password would be entered if the transaction were password protected.

### Authorizing commands

When the Command Authorization exit routine (DFSCCMD0) is included in IMS.SDFSRESL, it provides a command authorization check. The Command Authorization exit routine can work in conjunction with RACF, or it can work independently, without RACF .

The Command Authorization exit routine is called for each IMS command. Default security or RACF is called first to perform the authorization. The return code is passed to the Command Authorization exit routine. DFSCCMD0 performs a final verification and determines the success or failure of the command authorization. DFSCCMD0 can also be used alone to perform the verification. If you want to establish a command authorization level more discrete than that provided by RACF , you can examine DFSCCMD0's input command buffer. DFSCCMD0's input command buffer contains the complete command stream.

If DFSCCMD0 exists in IMS.SDFSRESL, DFSCCMD0 is called for all device types including static, ETO, and LU 6.2.

IMS provides the RVFY= parameter in the IMS procedure for customers who want to force reverification that the operator who signed on to a terminal is the same operator who is now entering a command or transaction. This reverification is done with RACF by including the word 'REVERIFY' in the APPLDATA field of the command or transaction profile. For example:

```
RDEFINE Cxxx tran-name UACC(NONE) APPLDATA('REVERIFY')
```

Each time the user enters this transaction code, the RACF password must be entered where an IMS password would be entered if the transaction were password protected.

### Password protecting the /LOCK, /UNLOCK, and /SET commands

For the /LOCK, /UNLOCK, and /SET commands you can implement password security to protect LTERMs, databases, programs, and transactions.



Perform the following tasks to check the password and enforce password security using RACF or an exit routine:

1. Define the resource to be protected by using the appropriate RACF security class: LIMS, PIMS, IIMS, or TIMS.
2. Define the user IDs that you want to have access to the resource in the RACF security class that defines the resource. The user IDs must be the same user IDs that are used to sign on to the terminal from which the /LOCK, /UNLOCK, or /SET command is issued.
3. In the DFSDCxxx PROCLIB member, specify LOCKSEC=Y.
4. In a DBCTL only environment, the /LOCK command cannot be secured with password protection.

## Using RACF to protect physical terminals

RACF offers a terminal-user security function that ranges from no security for a particular terminal to permitting a certain predefined list of users access through a physical control point. A terminal-user profile can be created for every PTERM in the IMS.

The following table is an example of a terminal-user profile.

*Table 23. Example terminal-user profile.*

| User   | Physical terminal |        |        |        |        |
|--------|-------------------|--------|--------|--------|--------|
|        | PTERMA            | PTERMB | PTERMC | LTERMD | LTERME |
| USER 1 | X                 |        |        |        | X      |
| USER 2 | X                 | X      |        |        |        |
| USER 3 |                   |        |        | X      | X      |
| USER 4 |                   |        |        |        | X      |
| USER 5 |                   | X      |        |        | X      |

Users defined through ETO can use signon verification to gain access to IMS transactions or commands. Dynamic terminal security must be defined through a security product such as RACF or through exit routines.

### Signon password with RACF

If RACF is used to implement signon verification security for a terminal, a check of the password is made with entry of the /SIGN command, by either the Signon Verification exit routine or by RACF. The user ID entered with the /SIGN ON command must be accompanied by a user password and, optionally, by other signon data. If the RACF reverification option has been specified, the password is saved so it can be compared with the password reentered with the transaction code.

### RACF password protection

RACF passwords are defined and maintained by the user. After the RACF resource class is initialized with a password, the user can change its value. If signon verification is provided by an exit routine and not by RACF, the table of user IDs and passwords must be changed by another bind into the IMS nucleus. However, for ETO signon verification, the table of user IDs and passwords does not need to be changed by another bind into the IMS nucleus. If the table is loaded by exit routine DFSINTX0 or if it is part of exit routine DFSSGNX0, IMS does not need to

be restarted in order to have the table refreshed.

**Related reference:**

 /SIGN command (Commands)

## Security considerations for the master terminal

The security of access from the master terminal is critical. Because the MTO can modify all security profiles during normal operations, you should consider protecting the terminal with a second level of control.

Signon verification security provides this capability. The primary question is how much capability to modify security should be given to this second level of control.

Default security does not and cannot prevent modifying the system's security profiles through the master terminal; however, you might want to restrict some commands from being entered from the MTO.

You can use the DFSCCMD0 exit routine to limit the commands that can be entered from the MTO.

Use the DFSCCMD0 exit routine to specify the commands that you want the MTO to be able to enter. The /ASSIGN, /CHANGE, and /DELETE commands are likely candidates to protect. For a cold start or warm start, the MTO can control the following security:

- Signon verification security
- Transaction authorization
- Transaction command security
- Command authorization

You can control these authorizations by making the specifications shown in the following table.

*Table 24. SECURITY macro keywords and initialization EXEC parameters*

| Enforced security option  | SECURITY macro keyword and parameter values | Initialization EXEC parameter |
|---------------------------|---------------------------------------------|-------------------------------|
| Signon verification       | SECLVL FORCSIGN                             | SGN                           |
| Transaction authorization | SECLVL FORCTRAN                             | TRN                           |
| Terminal security         | TERMNL FORCE                                | SGN                           |
| Command security          | TYPE RACFCOM                                | RCF                           |

## Security for AO application programs

Automated operator (AO) application programs can issue a subset of IMS operator commands. Because an operator command can compromise a security profile, you should prevent AO application programs from issuing sensitive commands that might modify IMS resources.

There are two types of AO application programs: type-1, which uses the CMD call to issue commands in DB/DC and DCCTL environments, and type-2, which uses the ICMD call to issue commands in DB/DC, DCCTL, and DBCTL environments. Your options for implementing security differ slightly between commands issued using CMD calls and commands issued using ICMD calls.

For both CMD and ICMD calls you must specify which security facilities will enforce AO security and then tailor the system definition macros and startup procedures to work with your security facility choices.

**Related reading:** For more complete information about AO applications, see *IMS Version 11 Operations and Automation*.

### **Type-1 AO command security**

You can use either RACF or another external security product, the Command Authorization exit routine (DFSCCMD0), or both of the security facilities to secure commands that are issued by type-1 AO application programs using the CMD call.

Make your specifications for security facilities using the AOI1= keyword in the IMS and DCC startup procedures.

Use the AOI1= execution parameter in the startup procedure to specify the security facilities that implement security for type-1 AO application programs. Specifications made using AOI1= override the specifications made in the TRANSACT and SECURITY macros.

You can change the AOI1= value each time you initialize IMS because the AOI1= specification is not included in checkpoint records.

**Related reading:** For detailed information about the AOI1= parameters and their relationship to specifications made using the AOI= keyword in the system definition macros, see *IMS Version 11 System Definition*.

### **Defining the transactions that issue type-1 AO commands:**

You need to define the transactions that are allowed to issue type-1 AO commands. You can do this using the AOI= parameter in the TRANSACT macro.

The AOI= parameter also allows you to specify the user ID used for security checking when you use RACF, DFSCCMD0, or both as your security facility. You can use any of the following identifiers as a user ID:

- The user ID of the user signed on at the terminal issuing the transaction
- The transaction code
- The first three letters of the commands name

**Related reading:** For more information on the AOI= parameter, see *IMS Version 11 System Definition*.

### **Type-2 AO command security**

You can use either RACF or another external security product, the Command Authorization exit routine (DFSCCMD0), or both of the security facilities to secure AO commands issued by type-2 AO application programs through the ICMD call.

Make your specifications for security facilities by using the AOIS= parameter in the IMS and DCC startup procedures.

Use the AOIS= execution parameter in the startup procedure to specify which security facilities to use for security checking of type-2 AO commands.

The AOIS= execution parameter allows you to specify one of the following options:

- RACF performs security checking.

- DFSCCMD0 performs security checking.
- RACF and DFSCCMD0 both perform security checking.
- Disable all ICMD calls from applications.
- Disable all security checking for ICMD calls.

You can change the AOIS= value each time you initialize IMS, because the AOIS= specification is not included in checkpoint records.

**Related reading:** For complete information on specifying the AOIS= keyword, see *IMS Version 11 System Definition*.

### **Defining the user ID for transactions that issue type-2 AO commands:**

Use the AOI= keyword on the TRANSACT macro to specify what RACF and DFSCCMD0 use as a user ID for security checking.

You can use any of the following identifiers as a user ID:

- The user ID of the user signed on at the terminal issuing the transaction
- The transaction code
- The first three letters of the commands name

In certain circumstances, the user ID that is used to issue a transaction might not be available when security checking is performed. In these cases, IMS substitutes other identifiers as the user ID. For information on when IMS substitutes identifiers and what IMS uses as a substitute user ID, see “User ID substitutions for AO application programs” on page 318.

Unlike type-1 AO commands, you cannot disable the ability of a transaction to issue type-2 AO commands. AOI=NO applies only to type-1 AO commands and does not prevent transactions from issuing type-2 commands. For type-2 AO commands, AOI=NO is the same as AOI=YES.

### **AO command security and system definition**

Make specifications that affect AO command security in the TRANSACT macro or with initialization EXEC parameters.

The AOI= keyword in the TRANSACT macro serves the following two functions:

1. You can specify which transactions can issue type-1 AO commands
2. If you use RACF or the Command Authorization exit routine (DFSCCMD0) for security checking, you can specify what the transactions provide as a user ID when invoking type-1 and type-2 AO application programs

For type-1 AO commands, the specifications that you make in the TRANSACT macro can be overridden using the AOI1= keyword in the startup procedure.

AOI=NO is ignored by type-2 AO application programs that use ICMD calls. You cannot disable the ability of a transaction to issue commands through type-2 AO application programs.

For both type-1 and type-2 AO application programs, you can specify the following user IDs for RACF and DFSCCMD0 to use for AO command authorization:

- The user ID that is signed on to the terminal that is issuing the transaction. Specify this with AOI=YES. Because a user ID is not always available for

security checking, IMS can substitute an LTERM name, PSB name, or other identifier as the user ID. For more information, see “User ID substitutions for AO application programs.”

- The transaction code. Specify this with AOI=TRAN.
- The first three letters of the command name. Specify this with AOI=CMD.

To increase your options when implementing AO command security, you can define any one of the three elements involved in security checking—the user ID, the transaction, or the command—as the user to be authorized.

For AO security, commands issued through CMD and ICMD calls are typically defined as a resource and either the signed-on user or the transaction code is defined as the user. However, you can instead choose to define the command name as the user to be authorized and the transaction code as the resource to be protected. Changing which element is defined as the user might be beneficial when, for example, you have so many different transactions that it becomes impractical to add them all to the RACF security class for a command defined as a resource.

**Related reading:** For detailed information about the TRANSACT macro and the AOI= keyword parameter, see *IMS Version 11 System Definition*.

## User ID substitutions for AO application programs

In certain circumstances, IMS substitutes other identifiers, such as an LTERM name or program name, as the user ID. Unless you have anticipated this and defined substitute user IDs to RACF, RACF rejects the CMD or ICMD call.

The following examples show how IMS might use user IDs in different circumstances in the AO application program's execution environments. These examples are based on a specification of AOI=YES in the TRANSACT macro.

### MPP or IFP

If a message GU call has completed, the user ID is used to determine if the user can issue commands using ICMD. The user ID is of a signed-on terminal or the LTERM name of the signed-off terminal where the transaction is issued. If GU is not issued, PSB name is used.

**BMP** If a message GU call has completed, the user ID is used to determine if the user can issue commands using ICMD. The user ID is of a signed-on terminal or the LTERM name of the signed-off terminal where the transaction is issued. If GU is not issued or if the BMP is non-message driven, the value of the USER parameter specified on the JCL JOB statement is used. If the USER parameter is not specified, a user ID of 0000000 is used.

### DRA THREAD

The security token that is passed in the PAPL for a schedule request is used to determine whether the user can issue commands using ICMD.

BMP regions are initiated by a /START command or by a JCL job. If you use a /START command, you should use RACF to protect the library that contains the cataloged procedures. If you use JCL, you can use RACF resource access security (RAS). You control access to the BMP region by defining a unique user ID to RACF on the JOB statement. If you are using RAS, enter the user ID in the RACF security class.

When AO application programs issue a command through the CMD or ICMD call, passwords are not used.

**Related reference:**

 IMS type-1 commands supported from an AO application (Commands)

## Implementing AO command security with RACF

RACF implements security based on user IDs and resource profiles stored in RACF security classes. RACF allows a user access to a resource only if the user's user ID is associated with the resource in the appropriate RACF security class.

For both type-1 and type-2 AO application program security, you must coordinate the definitions of the user IDs and resources that you make in the RACF resource class with your specifications in the AOI= keyword parameter in the TRANSACT macro.

You might need to define substitute user IDs depending on the type of dependent region in which your application programs are executing and other circumstances. For a description of those circumstances, see “User ID substitutions for AO application programs” on page 318.

## Example RACF definitions for AO application program security

When using RACF for security for both type-1 and type-2 AO application programs, you can use the following examples to help code your RACF definitions. The examples define the user IMSUSER1 and the transaction APOL13. The RACF definitions you use differ depending on the specifications that you make by using the AOI= keyword on the TRANSACT macro.

In the following example, where AOI=YES, the command /STOP is the protected resource and the user IMSUSER1 is authorized for the command. IMSUSER1 signs on to an IMS terminal and enters transaction APOL13. The transaction issues a CMD or ICMD call to issue the /STOP command (For type-2 AO security, the specification of AOI=YES is not required).

```
ADDUSER IMSUSER1 PASSWORD(IUSER1PW) DFLTGRP(SYS1)
RDEFINE CIMS (STO) UACC(NONE)
PERMIT STO CLASS(CIMS) ID(IMSUSER1) ACCESS(UPDATE)
```

In the following example, where AOI=TRAN, the command /STOP is the protected resource and the transaction APOL13 is authorized for the command. Any user can enter transaction APOL13. The transaction issues a CMD or ICMD call to issue the /STOP command. The APOL13 transaction is authorized to issue the /STOP command.

```
ADDUSER APOL13 NOPASSWORD DFLTGRP(SYS1)
RDEFINE CIMS (STO) UACC(NONE)
PERMIT STO CLASS(CIMS) ID(APOL13) ACCESS(UPDATE)
```

In the following example, where AOI=CMD, the transaction APOL13 is the protected resource and the user IMSUSER1 and the command /STOP are authorized for that transaction. Authorization is checked twice: when IMSUSER1 enters the transaction at the terminal and when the transaction issues a CMD or ICMD call to issue the /STOP command.

```
ADDUSER IMSUSER1 PASSWORD(IUSER1PW) DFLTGRP(SYS1)
ADDUSER STO NOPASSWORD DFLTGRP(SYS1)
RDEFINE TIMS (APOL13) UACC(NONE)
PERMIT APOL13 CLASS(TIMES) ID(IMSUSER1) ACCESS(UPDATE)
PERMIT APOL13 CLASS(TIMES) ID(STO) ACCESS(UPDATE)
```



**Related reading:**

- For more information about defining users for AO application program security purposes, see “AO command security and system definition” on page 317.
- For more information about programming RACF security classes, see *z/OS Security Server RACF System Programmer's Guide*.
- For more information about RACF commands, see *z/OS Security Server RACF Command Language Reference*.

**Implementing AO command security with the command authorization exit routine**

The Command Authorization exit routine (DFSCCMD0) is called during AO command processing to perform command authorization checking. DFSCCMD0 enables you to secure commands that are issued through either the CMD or ICMD call at the command verb, keyword, or resource name level.

DFSCCMD0 must be included in IMS.SDFSRESL.

The parameter list for DFSCCMD0 identifies:

- Who issued the command:
  - Terminal
  - LU 6.2 application
  - ICMD call, where a user ID is used for command authorization
  - ICMD call, where a PSB name is used for command authorization
- If RACF (or equivalent) was called:
  - SAF (System Authorization Facility) return code
  - RACF return code
  - RACF reason code
- The security code:
  - X'80000000' RACF was not called (AOIS=C).
  - X'00000000' User is authorized to RACF to issue command.
  - X'00000004' RACF is not available.
  - X'00000008' User is not defined to RACF.
  - X'0000000C' Command is not protected by RACF.
  - X'00000010' User is not authorized to issue command.

**Related reading:** For more information about the Command Authorization exit routine (DFSCCMD0), see *IMS Version 11 Exit Routines*.

## Security for Time-Controlled Operations

There are two ways that you can implement security for Time-Controlled Operations (TCO): by restricting the users (LTERMs) that can load a TCO script, and by restricting the IMS commands that a loaded TCO script can issue.

To restrict the LTERMs that can load a TCO script, modify the TCO CNT Edit exit routine (DFSTCNT0) without changing its name. IMS always calls DFSTCNT0 when a TCO script is loaded and, by default, DFSTCNT0 does not restrict the LTERMs.

To restrict the IMS commands that a TCO script can issue, specify A or S on the RCF execution parameter and Y on the TCORACF execution parameter in your startup procedure. IMS calls RACF and, if it exists, the Command Authorization

exit routine (DFSCCMD0). Finally, you must include /SIGN ON and /SIGN OFF commands at the beginning and end of the script, respectively. The user ID that is signed on at the beginning of the script is checked for authorization to the commands that the script issues.

**Related reading:**

- For more information about the RCF and TCORACF execution parameters, see “Procedures used in IMS environments” in *IMS Version 11 System Definition*.
- For more information about the DFSCCMD0 and DFSTCNT0 exit routines, see *IMS Version 11 Exit Routines*.
- For introductory information about TCO, see *IMS Version 11 Operations and Automation*.

## Security for Fast Path application programs

When you design security protection for Fast Path application programs or for DL/I programs that access Fast Path databases, consider protecting the processing in a dependent region and the network of Fast Path terminals.

Protect the dependent region and terminals by:

- Protecting the processing in a dependent region by securing the PSBs of the Fast Path application programs that are scheduled in the dependent region. You can do this by using RAS security and the RACF IIMS security class.
- Protecting the network of Fast Path terminals by using signon verification and transaction authorization.

## Security for JMP application programs

Application programs that run in JMP regions need to access UNIX System Services. Therefore, when running a JMP application, you must use RACF (or an equivalent product) and IMS signon verification security.

To use RACF with JMP regions, define the user IDs, which must have OMVS segments, to RACF.

**Related reading:** For information about defining user IDs to RACF, see *z/OS UNIX System Services Planning*.

## Security and CPI-C driven application programs

You can secure any PSB specified on an APSB call from a CPI-C driven application program using the z/OS System Authorization Facility (SAF).

After APSB SAF is security-enabled, IMS calls SAF to secure the PSB specified on an APSB call against the AIMS or Axxxxxxx general resource class (where xxxxxx is the value specified on the RCLASS= parameter of the IMS SECURITY macro or the RCLASS initialization EXEC parameter) based on the USERID of the user associated with the CPI-C application. Therefore, you must define the PSBs that you want protected by RACF (or your installation exit) to the AIMS or Axxxxxxx resource class. Because the AIMS resource class can contain PSBs, all PSB names specified in the AIMS resource class should be unique. In addition, you must specify RCLASS=IMS|xxxxxxx IMS system definition time.

When a CPI-C driven application program makes an APSB call, IMS bypasses APSB SAF security if the PSB is not defined to the AIMS resource class or if the AIMS resource class is not active. If IMS bypasses APSB SAF security, IMS



attempts to authorize the PSB by using RAS security. If RAS security is active, IMS assumes the CPI-C driven application program has the authority to use the PSB.

**Related reading:** For more information on RACF and IMS, search for IMS in the *z/OS Security Server RACF Security Administrator's Guide*.

## Security for ODBA application programs

You can secure any PSB specified on an APSB call from an ODBA application program using the z/OS System Authorization Facility (SAF).

After APSB SAF is security-enabled, IMS calls SAF to secure the PSB specified on an APSB call against the AIMS or Axxxxxxx general resource class, based on the user associated with the ODBA application at the task level (TCB) or the address space level (ASCB/ASXB) respectively. Therefore, you must define the PSBs that you want RACF (or your installation exit) to protect to the AIMS or Axxxxxxx resource class.

Since the AIMS resource class can contain PSBs, all PSB names specified in the AIMS resource class should be unique.

You must also specify RCLASS=IMS|xxxxxxx on the IMS SECURITY macro or with an initialization EXEC parameter during IMS system definition.

**Related reading:** For more information about RACF and IMS, search for “IMS” in the *z/OS Security Server RACF Security Administrator's Guide*.

## Use of the RACF data space

When the RACF (or an equivalent product) data space is supported (RACF 2.1 or later), IMS loads the RACF profiles for the IMS commands and transactions into that data space instead of the IMS control region.

Use of the RACF data space invalidates the IMS online change support for RACF with the /MODIFY command. The IMS online change support is still valid, though, when the RACF data space is not being used.

The message DFS3432 RACF PARAMETER INVALID IF RACF DATA SPACE USED is issued if the RACF parameter is used on the /MODIFY PREPARE command when the RACF data space is being used. You can use the RACF command SETROPTS RACLIST (*classname*) REFRESH to refresh the RACF resource profiles in the RACF data space without requiring the IMS applications to suspend work.

## Security in MSC and shared-queues environments

In Multiple Systems Coupling (MSC) and shared-queues environments that are made up of multiple IMS systems, each IMS system can receive transactions from terminals and from the other IMS systems in the same environment. The transactions from the other IMS systems require you to take additional security measures.

In an MSC network, each IMS system can be viewed as a local IMS system or a remote IMS system. A local IMS system receives transactions from a terminal. A remote IMS system receives transactions from a local IMS system through an MSC link. Each IMS system in an MSC environment can be local for one type of transaction and remote for another.

Similarly, in a shared-queues environment, each IMS system can function as a front-end IMS system, a back-end IMS system, or both. A front-end IMS system receives transactions from terminals. A back-end IMS system receives transactions from a front-end IMS system through a shared queue. Each IMS system in a shared-queues environment can function as both a front-end IMS system and a back-end IMS system.

In MSC and shared-queues environments, you must implement security separately for each IMS system. The security controls of each IMS system are independent of those of the other IMS systems. Each IMS system is unaware of the security checking that the other IMS systems perform.

**Related reading:** For more information about MSC and shared queues, see *IMS Version 11 Communications and Connections*.

### **Local MSC and front-end shared-queues security**

Although there are no special security specifications that you need to make for local MSC and front-end shared-queues IMS systems, if the IMS system might also function as a remote MSC or back-end shared-queues IMS system, be sure to implement security appropriately. Otherwise, implement security for local and front-end IMS systems as you would for a standalone IMS system in a DB/DC or DCCTL environment.

### **Remote MSC and back-end shared-queues security**

For remote MSC and back-end shared queues IMS systems, you can use RACF for security. CHNG and AUTH calls and IMS conversational deferred program switches (that occur in the same IMS as the inputting terminal) perform an authorization check to determine whether the user who entered the transaction is authorized to use the IMS resource.

To perform a RACF authorization from the dependent region, a security environment first must be established. If the dependent region is part of the same IMS as the inputting terminal, and the user who entered the transaction is still signed on, then the security environment that is created in the IMS control region at signon is used for the authorization call. The security environment that is created at signon is not available under the following conditions:

- The dependent region is part of the same IMS as the inputting terminal, but the user has signed off.
- The dependent region is part of another IMS, connected to the IMS with the inputting terminal by an MSC link.
- The dependent region is part of another IMS in a sysplex with IMS shared message queues support.

In these conditions, the security environment must be dynamically created in order to perform the RACF authorization check. Dynamic creation of this security environment increases the time required to process a CHNG or AUTH call. The dynamically created security environment is kept until IMS is done with the message (until the next sync point or GET UNIQUE). Use the Build Security Environment exit routine (DFSBSEX0) to control when IMS performs the dynamic creation of the security environment. to control when IMS performs the dynamic creation of the security environment.

**Related reading:** For an explanation of how to use the Build Security Environment exit routine (DFSBSEX0), see *IMS Version 11 Exit Routines*.

For input from APPC or OTMA, if security is defined as FULL, the security environment has already been created before any CHNG or AUTH call. If APPC/OTMA security is defined as NONE, RACF is not called. If APPC/OTMA security is defined as CHECK, and a CHNG or AUTH call is made, a dynamic security environment has to be created.

When the following IMS exit routines are called because of an application program CHNG or AUTH call, the address of the CTB is zero when the call is made from an IMS dependent region that is not part of the same IMS as the inputting terminal:

- Command Authorization exit routine (DFSCCMD0)
- Transaction Authorization exit routine (DFSCTRNO)
- Security Reverification exit routine (DFSCTSE0)

Resource Access Security (RAS) security for LTERMs in a shared-queues environment is supported only for LTERMs that are statically defined to that IMS back end.

### **Security for transactions that are received from an MSC link**

You have several options when specifying security for transactions that are received by an IMS system from an MSC link. By default, RACF and exit routines check security for direct-routed transactions and checks security for non-direct-routed transactions received from an MSC link.

Or, if your installation does not require MSC transaction security, you can disable all security checking for transactions received from an MSC link.

If you specify the use of RACF or exit routines to check security for any transactions received on an MSC link, you also need to specify which user ID these security facilities will use.

You can make all of the security specifications that are discussed in this topic by using the positional parameters of the MSCSEC= startup parameter in the DFSDCxxx PROCLIB member. The specifications made in MSCSEC= define the default security that your IMS system performs when it receives transactions from an MSC link.

The security specifications that you make with the MSCSEC= startup parameter are local in scope and apply only to the IMS system in which you make them.

**Related reading:** For more information about the DFSDCxxx PROCLIB member, see *IMS Version 11 System Definition*.

### **Specifying security for MSC transaction types:**

The MSCSEC= startup parameter, which is found in the DFSDCxxx PROCLIB member, includes two positional parameters. The first positional parameter allows you to specify—based on the MSC transaction type—if and when a receiving IMS calls RACF and exit routines to check security for a transaction.

There are two types of transactions that can be received on an MSC link for which you might perform security checking: direct-routed transactions and

non-direct-routed transactions. You specify which type of transactions that you want RACF and exit routines to check security for by using the first positional parameter of MSCSEC=. The options are:

**LRDIRECT**

RACF and exit routines check security on direct-routed transactions and no security checking is performed for non-direct-routed transactions. This is the default.

**LRNONDR**

RACF and exit routines check security on non-direct-routed transactions and no security checking is performed for direct-routed transactions.

**LRALL**

RACF and exit routines check both types of transactions.

**LRNONE**

The IMS system does not request any security checking for either type of transaction.

When you use RACF and exit routines to check security for transactions that are received from an MSC link, the receiving IMS system calls RACF and DFSCTRN0 after calling the TM and MSC Message Routing and Control user exit routine (DFSMSCE0).

**Specifying user IDs for MSC security:**

If you use RACF, exit routines, or both to perform MSC security checking, use the second positional parameter of the MSCSEC= startup parameter in the DFSDCxxx PROCLIB member to specify the user ID on which to base security checking.

The parameter options are:

**CTL**

Specifies that RACF and exit routines use the user ID of the receiving IMS control region. This is the default.

**MSN**

Specifies that RACF and exit routines use the MSNAME as the user ID.

**USR**

Specifies that RACF and exit routines use the user ID of the inputting terminal.

**Note:** User IDs that are defined in RACF resource classes must be unique. When you define user IDs for MSC security, watch for conflicts with existing user IDs.

**MSC security and the TM and MSC message routing and control exit routine:**

The TM and MSC Message Routing and Control exit routine (DFSMSCE0) is concerned primarily with the routing of messages for TM and MSC; however, you can also use this exit to specify which user ID will be used for security checking.

The specifications you can make with DFSMSCE0 regarding user IDs are the same as those specified by the second positional parameter of the MSCSEC= keyword in the DFSDCxxx PROCLIB member.

When IMS calls DFSMSCE0 for MSC link-receive routing, IMS passes to DFSMSCE0 the second security specification defined in the DFSDCxxx MSCSEC= parameter. If the security specification in DFSMSCE0 differs from the one received

from IMS for a given transaction, DFSMSCE0 can override the specification received from IMS for that transaction. DFSMSCE0 can also reroute or cancel the transaction.

IMS calls DFSMSCE0 through one of four entry points when transactions are received from an MSC link, but only the following two entry points are used for security checking:

- LRTRAN for non-direct-routed transaction messages local to the receiving IMS system
- LRDIR for direct-routed transaction messages local to the receiving IMS system

**Related reading:** For more information about the TM and MSC Message Routing and Control exit routine (DFSMSCE0), see *IMS Version 11 Exit Routines*.

#### **MSC security and the Transaction Authorization exit routine:**

The Transaction Authorization exit routine (DFSCTRN0) works with the Security Reverification exit routine (DFSCTSE0) and the Signon/off Security exit routine (DFSCSGN0) to check an individual user ID for authority to use a transaction.

DFSCTRN0 can be used with or without RACF to verify that the user ID is authorized to run a transaction. If both the RACF option and DFSCTRN0 are selected in the IMS system definition, the exit routine is activated after RACF verifies the transaction. If the transaction request is rejected by RACF, the exit routine is not called. If the RACF option is not selected in the IMS system definition, this exit routine can be used to verify the user's authorization and the password, if required, for that transaction.

**Related reading:** For more information about the Transaction Authorization exit routine (DFSCTRN0), see *IMS Version 11 Exit Routines*.

## **Security for APPC/IMS**

APPC/MVS does not verify user authority to access transaction codes or specific IMS systems. To provide complete security verifications of a user's authority to execute transactions with LU 6.2 devices, you must use RACF.

APPC/MVS uses RACF resource class APPCTP for security. This holds a profile for every IMS transaction defined to RACF for transaction authorization verification. The Transaction Authorization exit routine (DFSCTRN0) is called for transactions, and the Command Authorization exit routine (DFSCCMD0) is called for commands.

### **Security checking with APPC local LUs**

You can use either the base logical unit (LU) or local LU for RACF security checking of outbound asynchronous messages. Use the APPCLU startup parameter on the DFSDCxxx PROCLIB member data set to specify whether you want the base or a local LU to be used for asynchronous message responses.

In an IMS Version 11 shared-queues environment, members that could not support local logical units (LUs) can tolerate local LUs; however, the resource information about those local LUs is lost. If an IMSplex member does not support local LUs, then the base LU is used to send asynchronous messages.

It is recommended that all members of an IMSplex migrate to IMS Version 11 before using the APPCLU startup parameter. See *IMS Version 11 System Definition* for information on the APPCLU startup parameter.

## Security for ETO terminals

ETO provides dynamic LTERM support. You can dynamically create and allocate local LTERMs to a terminal session based on user signon or the application ISRT process. An LTERM can be associated with a specific user ID instead of a physical terminal. By associating an LTERM with a user ID, you prevent the wrong user from receiving messages at a physical terminal.

**Related reading:** For additional information about ETO security, see “Planning a high-security environment with ETO” in *IMS Version 11 Communications and Connections*.

The Initialization exit routine (DFSINTX0), which is called before IMS loads ETO descriptors, can cause user-defined security information to be loaded and made available to security exit routines. DFSINTX0 can also be used to request that IMS reverify new passwords during signon.

After calling DFSINTX0, IMS searches for and uses the following security exit routines in order:

1. Logon exit routine (DFSLGNX0)

You can include the Logon exit routine (DFSLGNX0) in IMS.SDFSRESL when ETO is implemented. DFSLGNX0 is an installation exit routine. Among the other functions of DFSLGNX0, it enables you to do the following:

- Allow or disallow a logon attempt based on criteria that you specify.
- Create or modify the user data that you want IMS to pass to the Sign-on exit routine (DFSSGNX0).

2. Sign-on exit routine (DFSSGNX0)

You can include the Sign-on exit routine (DFSSGNX0) in IMS.SDFSRESL when ETO is implemented. DFSSGNX0 is a user exit routine. Among its other functions, it enables you to allow or disallow a signon attempt based on any criteria that you specify.

You can also use the DFSSGNX0 exit routine to dynamically create node user descriptors and to create a user structure name that is different from the user ID, using a suffix and the node name.

**Related reading:** For more information about DFSSGNX0, DFSLGNX0, DFSINTX0, see *IMS Version 11 Exit Routines*.

## Securing DB/DC and DCCTL dependent regions and their resources

It is important to protect dependent regions. There are two approaches to securing a dependent region: preventing the unauthorized scheduling of application programs in the dependent region and preventing the unauthorized use of resources by those programs after they have been scheduled.

The following list describes the different types of dependent regions and the resources that can be protected for each:

- For BMP regions:
  - Transaction codes



- PSBs
- Logical terminals or transaction codes that are specified by using the OUT= keyword of the IMSBATCH procedure
- For IFP regions: PSBs
- For JBP regions:
  - PSBs
  - Logical terminals or transaction codes that are specified by using the OUT= keyword of the DFSJBP procedure
- For JMP regions: transaction codes
- For MPP regions: transaction codes

IMS supports Resource Access Security (RAS) to protect dependent regions.

### **Securing dependent regions using resource access security**

Resource access security (RAS) prevents an application program that is running in a dependent region from using a resource (a transaction, PSB, or LTERM) unless it is authorized to do so. RAS does not restrict the scheduling of application programs in dependent regions.

The authority of an application program to access a resource that is protected by RAS is based on the user ID of the dependent region. The user ID of the dependent region must be authorized in the RACF security class profile for the resource that the application program is attempting to use.

You can specify RAS by using the TYPE= keyword parameter in the SECURITY system definition macro or by using the ISIS= execution parameter in the IMS, DBC, and DCC startup procedures.

#### **Related reading:**

- For more information about coding the SECURITY macro and its parameters, see “Defining the SECURITY macro for DB/DC and DCCTL environments” on page 331.
- For additional information about the IMS, DBC, or DCC startup procedures, see *IMS Version 11 System Definition*.

### **RACF resource classes for RAS security**

RAS uses the RACF resource classes to define the resources it protects and the user IDs of the dependent regions that can access those resources. The RACF resource classes that RAS uses include:

- IIMS and JIMS for PSBs and groups of PSBs
- LIMS and MIMS for LTERMs and groups of LTERMs
- TIMS and GIMS for transactions and groups of transactions

Some of these classes are predefined by RACF, but others are not. If you do not find the security class that you need among the classes RACF provides, you can use the RACF resource class macro ICHERCDE to create an installation-defined class descriptor table (CDT).

The RACF security classes contain the names of the resources that RAS protects and the user IDs that can use them. When an application program in a dependent region attempts to use a resource, RAS checks the class profile of the resource to

see if the user ID of the dependent region is listed as authorized for that resource. If the user ID is specified, RAS grants access to the application program. If the user ID is not specified, RAS denies access.

#### **Related reading:**

- For more information about working with RACF security, see:
  - “Preparing a RACF security plan” on page 332
  - *z/OS Security Server RACF Security Administrator’s Guide*
  - *z/OS Security Server RACF Macros and Interfaces*
- For more information about defining new classes with the ICHERCDE macro, see *z/OS Security Server RACF System Programmer’s Guide*.

### **Authorizing resource use in a dependent region**

For the IMS dependent region, you can use RACF (or an equivalent product) for security. CHNG and AUTH calls and IMS conversational deferred program switches that occur in the same IMS as the inputting terminal perform an authorization check to determine whether the user who entered the transaction is authorized to use the IMS resource.

To perform a RACF authorization from the dependent region, a security environment must be established. If the dependent region is part of the same IMS as the inputting terminal and the user who entered the transaction is still signed on, then the security environment that is created in the IMS control region at signon is used for the authorization call. The security environment that is created at signon is not available under the following conditions:

- The dependent region is part of the same IMS as the inputting terminal, but the user has signed off.
- The dependent region is part of another IMS, connected to the IMS with the inputting terminal by an MSC link.
- The dependent region is part of another IMS in a sysplex with IMS shared message queues support.

In these conditions, the security environment must be dynamically created to perform the RACF authorization check. Dynamic creation of this security environment increases the time required to process a CHNG or AUTH call. The dynamically created security environment is kept until IMS is done with the message (that is, until the next sync point or GET UNIQUE).

For input from APPC or OTMA, if security is defined as FULL, then the security environment is already created before any CHNG or AUTH call. If APPC or OTMA security is defined as CHECK and a CHNG or AUTH call is made, then a dynamic security environment has to be created.

The user ID that is used for creating the security environment for making the authorization call is based on the following environments and criteria:

#### **MPP or IFP**

If a message GU call has completed, then the security value from the input message is used to perform the authorization. The security value is the user ID of a signed-on terminal or the LTERM name of the signed-off terminal where the transaction is issued. If a GU has not been issued, then the PSB name is used.

#### **BMP**

If a message GU call has completed, then the security value from the input message is used to perform the authorization. The security value is the



user ID of a signed-on terminal or the LTERM name of the signed-off terminal where the transaction is issued. If a GU has not been issued or if the BMP is non-message driven, then the value of the USER= parameter that is specified on the JCL JOB statement is used. If the USER= parameter is not specified, then a user ID of 0000000 is used.

If no accessor environment element (ACEE) exists in the IMS control region and a dynamic security environment cannot be created dynamically, then a default security environment is used. If an IMS BMP has PARDLI=1 specified or an IMS system is specified with LSO=Y, then the default security environment is the environment of the IMS control region that is created with the user ID that is associated with the IMS control region. Otherwise, the default security environment is that of the IMS dependent region that is created with the user ID that is associated with the IMS dependent region.

When the following IMS exit routines are called because of an application program CHNG or AUTH call, the address of the CTB is zero if the call is made from an IMS dependent region that is not part of the same IMS as the inputting terminal:

- Command Authorization exit routine (DFSCCMD0)
- Transaction Authorization exit routine (DFSCTRN0)
- Security Reverification exit routine (DFSCTSE0)

## Security for OTMA

Security for OTMA is enforced by RACF or a similar security product.

You can provide OTMA security by using:

- RACF security levels
- System execution procedures, the transaction message prefixes, and the /SECURE OTMA command
- Protecting messages on asynchronous hold queues from unauthorized use

### Related tasks:

 OTMA security (Communications and Connections)

## Security for IMS Connect

IMS Connect includes a variety of options for implementing and modifying the security checking performed on messages as they arrive in IMS Connect and as they arrive at the data store.

IMS Connect provides two options for checking security within IMS Connect: you can configure IMS Connect to call RACF directly or you can have the IMS Connect user message exit routines call a Security user exit routine.

Additional security features provided by IMS Connect include:

- Password management support
- Trusted-user classification for messages arriving at the data store
- OTMA accessor environment element (ACEE) timeout specification support
- Secure Socket Layer (SSL) support

### Related concepts:

 IMS Connect security support (Communications and Connections)

---

## Activating IMS security for DB/DC and DCCTL environments

This topic provides guidance on activating your IMS security design and on using RACF and program exit routines.

### Defining the SECURITY macro for DB/DC and DCCTL environments

The SECURITY macro statement allows you to specify optional security features to be in effect during IMS execution, unless they are overridden during system initialization.

Using the SECURITY macro, you can define the IMS system to use the Resource Access Control Facility (RACF) licensed program (or equivalent), an exit routine, or both to perform the following security functions:

- Transaction authorization
- Command authorization
- Signon verification
- PSB authorization verification
- Application resource access authorization
- DBCTL access authorization

Authority to issue specific IMS commands can be restricted to certain transactions, and password and terminal security requirements can be imposed.

If the SECURITY macro is included, any values specified or accepted by default override options in the COMM or IMSGEN macro.

**Related reading:** For detailed guidance on how to specify all of the keywords in the SECURITY macro, see *IMS Version 11 System Definition*.

### Preparing security exit routines

Each of the security exit routines need to be prepared as part of authorization.

- Signon/off Security exit routine (DFSCSGN0)

The /SIGN ON/OFF Security exit routine must be coded by your installation as module DFSCSGN0. This exit routine should have access to a table of valid user IDs and their associated passwords and RACF PassTickets (if they are used). For addressability, the table should reside in module DFSCSGN0, the Transaction Authorization exit routine (DFSCTR0), or in the IMS nucleus. The exit routine should note each successful signon. When the /SIGN OFF command is executed, the exit routine should mark that user ID available for /SIGN ON. The exit routine can place information in the data portion of the user verification string for logging. (An address in a register points to the user verification string.)

- Transaction Authorization exit routine (DFSCTR0)

The Transaction Authorization exit routine (DFSCTR0) should have access to a table of valid user IDs, RACF PassTickets, passwords, and transactions associated with each valid user ID. For addressability, this table should reside in module DFSCSGN0, the /SIGN ON/OFF Security exit routine (DFSCSGN0), or in the IMS nucleus. If the table is in the nucleus, it can be shared by the Transaction Authorization exit routine and the Signon Verification exit routine.

If you use message edit routines, security is checked after the message is edited.

- Command Authorization exit routine (DFSCCMD0)

The Command Authorization exit routine (DFSCCMD0) should have access to a table of valid user IDs, passwords, and commands associated with each valid user ID. For addressability, this table should reside in module DFSCCMD0, the /SIGN ON/OFF Security exit routine (DFSCSGN0), or in the IMS nucleus. If the table is in the nucleus, it can be shared by the Command Authorization exit routine, the Transaction Authorization exit routine, and the Signon Verification exit routine.

- If you are using the Resource Access Security exit routine (DFSRS00)  
You can use the Resource Access Security exit routine (DFSRS00) to augment or refine the security functions that are provided by RAS. DFSRS00 is called after a call to RACF to directly authorize an IMS dependent region to use IMS resources (transactions, PSBs, and LTERMs) the dependent region attempts to access. You can specify DFSRS00 by using the TYPE= keyword in the SECURITY system definition macro or by using the ISIS execution parameter.

**Related tasks:**

➡ Using the Transaction Authorization exit routine (DFSCTRN0)  
(Communications and Connections)

**Related reference:**

- ➡ Resource Access Security exit routine (DFSRS00) (Exit Routines)
- ➡ Command Authorization exit routine (DFSCCMD0) (Exit Routines)
- ➡ Signon/off Security exit routine (DFSCSGN0) (Exit Routines)

## Preparing a RACF security plan

A security plan for RACF includes defining which resources need to be protected, specifying security options in system definition macros, and defining the resources that you want to protect to RACF.

To prepare a security plan that uses RACF:

1. Prepare a list of all of the IMS online resources to be protected, arranging them in groups to give an overview of the total resources covered.
2. Select the security facilities that protect the resource groups.
3. Design screen formats to include non-display fields for passwords in transactions and commands.
4. Specify your security options in the SECURITY, COMM, and MSGEN system definition macros.
5. Define the RACF resource class profiles to RACF.
6. Add users, groups, and data sets to RACF.
7. Define transactions and transaction groups to RACF.
8. Define databases, segments, fields, and other resources and resource groups to RACF.
9. Define commands and command groups to RACF.
10. Define extended resource protection sources (APPL).
11. Modify JCL procedures in IMS.PROCLIB.

IMS security uses a variety of RACF resource classes. These classes define individual resources or groups of resources, and they are divided into the following categories:

## APPC/IMS

The APPC resource classes used by APPC/IMS are not specific to IMS and include:

### APPCTP

Identifies transaction profiles for LU 6.2 transactions to RACF.

### APPCLU

Specifies the conversation security for a session.

### APPCPORT

Controls access to the system from a given LU (APPC port of entry).

**Related reading:** For additional information on APPC/IMS and the RACF resource classes used by APPC/IMS, see the APPC Transaction Security topic in *IMS Version 11 Communications and Connections*.

## Application

The application resource class, APPL, holds a profile of every subsystem that is defined to RACF. The application group resource class, AIMS, holds a profile for every APSB that is defined to RACF. The IMS system is defined in this class with the IMSID name (with the IMSCTRL macro) for system access authorization checking at signon.

## Command

The command resource class, CIMS, holds a profile for every command that is defined to RACF for command authorization checking. The command group resource class, DIMS, allows grouping of IMS commands that have a common access authority profile. Commands are defined for authorized user IDs.

## Database

The database resource class, PIMS, holds a profile for each database that is defined to RACF for authorization checking. The database group resource class, QIMS, allows grouping of database resources that have a common access authority profile.

**Field** The field resource class, FIMS, allows RACF authorization checking of fields within a database. The field resource group class, HIMS, allows grouping of common access fields for RACF authorization checking.

## LTERM

The LTERM resource class, LIMS, holds a profile for every IMS LTERM that is defined to RACF for LTERM authorization checking. The LTERM group resource class, MIMS, allows grouping of IMS LTERMs that have a common access authority profile.

**Other** This resource class, OIMS, and its group resource class, WIMS, are available for you to use for resources that do not fit into any other class. This calls can be used to interface with the AUTH call.

**PSB** The PSB resource class, IIMS, holds a profile for every IMS PSB that is defined to RACF for PSB authorization checking. The PSB group resource class, JIMS, allows grouping of IMS PSBs that have a common access authority profile.

## Segment

The segment resource class, SIMS, identifies individual segments to RACF . The segment group resource class, UIMS, allows grouping of segments with a common access authority profile for RACF authorization checking.

## Transaction

The transaction resource class, TIMS, holds a profile for every IMS transaction defined to RACF for transaction authorization checking. The transaction group resource class, GIMS, allows grouping of IMS transactions that have a common access authority profile.

Use the RCLASS= keyword in the SECURITY macro or the RCLASS= initialization EXEC parameter to specify the names of the resource classes that the IMS system will use.

The following table shows resource class assignments.

Table 25. Resource class assignments

| Resource class type                   | Resource class name                 |                                             |
|---------------------------------------|-------------------------------------|---------------------------------------------|
|                                       | RACF-defined name                   | User-defined name                           |
| APPC/IMS                              | APPCTP, APPCLU, APPCPOR, and others | These are not IMS-specific resource classes |
| Application resource class            | APPL                                | This is not an IMS-specific resource class  |
| Application group name resource class | AIMS                                | Axxxxxxx                                    |
| Command resource class                | CIMS                                | Cxxxxxxx                                    |
| Command group resource class          | DIMS                                | Dxxxxxxx                                    |
| Database resource class               | PIMS                                | Pxxxxxxx                                    |
| Database group resource class         | QIMS                                | Qxxxxxxx                                    |
| Field resource class                  | FIMS                                | Fxxxxxxx                                    |
| Field group resource class            | HIMS                                | Hxxxxxxx                                    |
| LTERM resource class                  | LIMS                                | Lxxxxxxx                                    |
| LTERM group resource class            | MIMS                                | Mxxxxxxx                                    |
| Other resource class                  | OIMS                                | Oxxxxxxx                                    |
| Other group resource class            | WIMS                                | Wxxxxxxx                                    |
| PSB resource class                    | IIMS                                | Ixxxxxxx                                    |
| PSB group resource class              | JIMS                                | Jxxxxxxx                                    |
| Resume TPIPE class                    | RIMS                                | Rxxxxxxx                                    |
| Segment resource class                | SIMS                                | Sxxxxxxx                                    |
| Segment group resource class          | UIMS                                | Uxxxxxxx                                    |
| Transaction resource class            | TIMS                                | Txxxxxxx                                    |
| Transaction group resource class      | GIMS                                | Gxxxxxxx                                    |

The RACF resource classes are defined in the RACF resource class descriptor table (CDT). Initially, the RACF-defined resource classes (shown in column 2 of Table 25) are predefined in the CDT. To add a resource class or to define a resource class with a user-defined name, you must use the RACF resource class macro ICHERCDE to create an installation-defined CDT.

**Requirement:** If you are going to allow mixed-case passwords (and you are using RACF for security), you must issue the SETROPTS,PASSWORD(MIXEDCASE) command.

**Related reading:**

- For more information about the ICHERCDE macro, see *z/OS Security Server RACF Macros and Interfaces*.
- For more information about updating the RACF resource class descriptor table, see *z/OS Security Server RACF System Programmer's Guide*.

**Related reference:**

 AUTH call (Application Programming APIs)

## Enabling and disabling APSB SAF security

You can enable APSB SAF security using one of the following methods.

- Specify RACF=FULL in the TP scheduler section of the CPI-C application's TP profile and issue the IMS command /SECURE APPC PROFILE. The command /SECURE APPC PROFILE enables APSB SAF Security only for the CPI-C applications that have RACF=FULL specified in the TP profile. The command disables APSB SAF Security for all other CPI-C applications.
- Issue the IMS command /SECURE APPC FULL to enable APSB SAF security for all CPI-C applications.

To disable APSB SAF security, issue the IMS command /SECURE APPC CHECK or /SECURE APPC NONE.

---

## Controlling security during system startup for DB/DC and DCCTL

The EXEC parameters in the IMS and DCC startup procedures provide a way to control the kind of security checking that is done during the current execution. The parameters act as switches for the different types of security that are specified in the system definition macros.

They also determine what flexibility the MTO has to override the choice of security checking. You must coordinate the setting of these parameters with both overall security design and operational procedures. The EXEC parameters for security are RCLASS, SECCNT, TRN, SGN, RCF, ISIS, ASOT, ALOT, AOI1, AOIS, and TCORACF.

The default values for the IMS and DCC procedures all specify no security. You must reset them to enable security.

The security functions and the EXEC parameters that you use to specify them are shown in the following table. Some of the EXEC parameters in the table override or replace related parameters specified in the system definition macros.

You must match the level of the security tables with the suffix identifier for the nucleus. Operational restrictions for the MTO are described in "Security considerations for the master terminal" on page 315.

**Related reading:** For detailed information on the JCL parameters and their parameter values, see *IMS Version 11 System Definition*.

Table 26. EXEC parameters to control IMS security

| Choice of security function                                        | EXEC parameter | Parameter value for security choice |                                                | Notes      |
|--------------------------------------------------------------------|----------------|-------------------------------------|------------------------------------------------|------------|
|                                                                    |                | Disable                             | Enable                                         |            |
| Identification of IMS to RACF as a resource class                  | RCLASS         | do not specify RCLASS               | specify a name on RCLASS and also specify RCF= |            |
| Number of security violations before MTO is notified               | SECCNT         | 0                                   | 1, 2, 3                                        |            |
| Transaction authorization                                          | TRN            | N                                   | Y, F, X, E                                     | 1, 2, 8    |
| Signon verification                                                | SGN            | N                                   | D, E, F, M, W, X, Y, Z, G,                     | 1, 2, 3, 8 |
| RACF security for transaction authorization or signon verification | RCF            | N                                   | A, B, C, R, S, T, Y                            | 4, 5, 8    |
| RAS dependent region security                                      | ISIS           | 0, 1, 2, N                          | A, C, R                                        | 6, 9       |
| Autosignoff                                                        | ASOT           | 0 or 1440                           | 10-1439                                        | 7          |
| Autologoff                                                         | ALOT           | 0 or 1440                           | 10-1439                                        |            |
| Security checking for CMD calls                                    | AOI1           | N                                   | A, C, R, S                                     | 8          |
| Security checking for ICMD calls                                   | AOIS           | S                                   | A, C, N, R                                     | 8          |
| RACF check of TCO-issued commands                                  | TCORACF        | N                                   | Y                                              | 8          |

**Notes:**

1. With value N, on the /NRESTART command, the MTO can optionally invoke checking.
2. With value Y, the security function is active unless overridden by the MTO.
3. Value M indicates multiple signons for a single user ID. Value Z is equivalent to Y + M; value G is equivalent to F + M.
4. The RACF licensed program is used in conjunction with Command Authorization, Transaction Authorization, or Signon Verification exit routines.
5. If a null value is specified, the choice is the default to that given in system definition.
6. The ISIS keyword parameter allows you to choose the type of dependent region security you want, RAS, and which security facilities, RACF , and exit routines, the security type will use.
7. On a terminal defined with ETO, when the last autologon user's last queue is completed, the autologon user immediately signs off without waiting for the autosignoff timeout interval.
8. Because this specification is not included in a checkpoint record, you can change its value each time IMS is initialized.
9. ISIS parameter values of 0, 1, or 2 are tolerated for compatibility. Internally, these values are converted to ISIS=N.



---

## Implementing security changes online in DB/DC and DCCTL environments

You can perform changes to your security definitions online when you use RACF for security enforcement.

To change RACF security definitions online, update the RACF database. Then issue the RACF command SETROPTS RACLIST(*classname*) REFRESH.

RACF refreshes the class that corresponds to *classname* and all classes that share the same POSIT value on their class descriptor table (CDT) entries.

**Related reading:** For more information on the SETROPTS command, see *z/OS Security Server RACF Command Language Reference*.

---

## Controlling security violations in DB/DC and DCCTL environments

Security violations are handled according to the installation's security administration guidelines.

IMS records the following security violation attempts on the IMS system log:

- Input message from an unauthorized terminal
- Password omitted when one is required
- Password incorrect for authorization
- Misspelled password
- Rejected signon
- Unauthorized DL/I command (CMD) call from application program

IMS rejects invalid input messages by sending a message to the terminal entering the message and logging the violation. The IMS system log provides an audit trail for investigation of possible security problems. The IMS system log security violation is identified as a X'10' log record type. You can use the File Select and Formatting Print utility to print the log.

You might want to have tighter security so that you are immediately notified about security violations. You can arrange for the master terminal to be immediately notified about security violations by having messages sent to it whenever the violations occur. To have the master terminal notified when violations occur, specify a non-zero value for the SECCNT keyword in the SECURITY macro or the SECCNT initialization EXEC parameter.

However, in a large network, misspelled passwords, transaction codes, and commands can cause an extremely large number of violations and violation notifications. You can reduce the number of notifications caused by operator errors, while still providing evidence of real attempts to avoid security safeguards, by specifying a notification threshold. When the number of violations from a single terminal equals the notification threshold value (as specified by the SECCNT value), the master terminal is notified.

Another method for recording security violations is available when RACF is installed. Each resource access violation creates a RACF type 80 record. You can use the RACF report writer to create reports based on these records.

**Related reading:**



- For more information on using the SECCNT keyword, see “Defining the SECURITY macro for DB/DC and DCCTL environments” on page 331.
- For more information about the SECCNT initialization EXEC parameter, see *IMS Version 11 System Definition*.
- For more information on the File Select and Formatting Print utility, see *IMS Version 11 System Utilities*.
- For more information on using the RACF report writer to format and print RACF records, see *z/OS Security Server RACF Auditor’s Guide*.

---

## Considering other access control methods

Use other access methods to augment your security software.

### Physical security

You should consider physical security measures that support your system security.

These measures include:

- Controlled access to and from the computer area
- Authorization of DP operations and non-operations personnel in certain terminal areas
- Separately controlled areas for media such as tapes, disks, cards, or files
- Control of computer forms and printed output

Physical security needs are likely to be dynamic and merit periodic review and adjustment.

### Using display bypass and password masking in DB/DC and DCCTL

IMS does not provide a software function to blank out or obliterate passwords from the terminal device display media after they are accepted. However, Message Format Service (MFS) facilities enable users to define fields with a non-display attribute (for 3270 display devices). IMS removes passwords from messages prior to recording them on the log.

If you plan to use passwords as part of transaction and command entry, you should design screen formats to incorporate non-display fields. This protection is especially important for the /SIGN command. The DFS3649 signon required message has non-display fields built into it for entering passwords on ACF/VTAM display terminals.

Most key-driven terminals have a feature (called the bypass feature) that permits characters to be entered without displaying them. Ordinarily, a terminal with this feature is operated continuously either in display or bypass mode. If passwords are to be masked to support security requirements, this feature is a necessity.

The bypass feature can be used operationally for establishing standards of protection for not only passwords, but also command verbs, commands, transaction codes, and text.

If passwords or other sensitive data must be altered in the IMS log, IMS provides a user exit for this purpose. For information about the Log edit exit routine (DFSFLGE0), see *IMS Version 11 Exit Routines*.

## Protecting your resources

You can protect IMS system libraries and data sets, as well as VSAM, OSAM, and Fast Path databases, in both the online (DB/DC, DBCTL, and DCCTL) and batch environments.

### IMS system libraries and system data sets

You can use RACF to protect IMS system libraries and system data sets. IMS invokes RACF to determine whether the user ID associated with the system address space (control region, DLISAS, or batch) attempting to open the resource has the necessary access authorization. Actually, when RACF authorizes access, it associates a user ID with the started procedure name (IMS or DLISAS procedure) through a started task table. If you start IMS with JCL, the RACF user ID can be on the job card along with its password.

If the user ID does not have the authorization, access is denied. The basic rule is “Whoever has the DD card must have the authority.”

**Related reading:** For more information on this process, see *z/OS Security Server RACF Security Administrator's Guide*.

### IMS procedure

If the IMS procedure is associated with a RACF user ID (with sufficient authority), the IMS control region can open a RACF-protected data set. If an association does not exist, the IMS control region is not allowed to open a RACF-protected data set that does not allow universal access for the requested authority level.

### DLISAS procedure

If the DLISAS procedure is associated with a RACF user ID, it overrides the RACF user ID for the IMS procedure. If an association does not exist, the RACF user ID associated with the IMS procedure is used for RACF access checking.

## Protecting databases

You can protect your VSAM and OSAM full-function databases, as well as your Fast Path DEDBs by using segment-level sensitivity, field-level sensitivity, or RACF.

**Segment- and field-level sensitivity:** Through centralized control over the content of database definitions, program specification blocks, and the libraries in which they reside, an effective scheme of protection attributes can be assigned to data. Note, however, that for database protection through PSBs to be totally effective, you should also protect the PSB library and the application program library (to safeguard the code that accesses the databases).

### Segment-level sensitivity

If you are not using field-level sensitivity, the smallest unit of data that can be protected is the segment. The basic actions that can be authorized are:

**None** No access to segment type.

**Read** Segment type can only be retrieved.

One or more of the following additional actions combined with read can be authorized:

**Add** New occurrences of segment type can be inserted.

**Update**

An existing occurrence of a segment type can be replaced.

**Delete** An existing occurrence of a segment type can be deleted.

The way the PCB and the parameter values for the PROCOPT keyword are specified controls the authorization. Although access authorization is declared at the program level, enforcement of the authorization can be made to appear at the transaction code or individual hierarchic level of a database. If only one transaction code is associated with a particular program, then the access authorization is effective at the transaction level. By using SENSEG statements in the PSB and key sensitivity as a processing option for higher-level segments, masking can be effective at the individual hierarchic level.

**Field-level sensitivity**

Field-level sensitivity can provide another kind of database security. Database descriptions (DBDs) and PSBs can be coded to permit access to a required subset of fields within a segment. Field-level sensitivity can also be used to control the replace function at the field level to help ensure database integrity.

**Related reading:**

- For more information on security at the database level, see *IMS Version 11 Database Administration*.
- For information about generating DBDs and PSBs, see “Application Control Blocks Maintenance utility” and “Program Specification Block (PSB) generation utility” in *IMS Version 11 System Utilities*.

**Protecting databases with RACF**

You use the RACF user ID of the DLISAS or control region started procedure, depending on the environment you are executing. If you start IMS with JCL, the RACF user ID can be on the job card along with its password.

**VSAM full-function database**

In an online environment, if a RACF user ID is associated with the DLISAS started procedure, that ID is used for access checking. If a RACF user ID is not associated with the DLISAS started procedure, the control region RACF user ID is utilized. (In the batch environment, the user ID of the batch job is employed.) Access authority of “CONTROL” is required. The database must be defined as ICFCATALOG. Use of the older “VSAM” catalog type is restricted.

**OSAM full-function database**

In an online environment, the RACF user ID of DLISAS is used; in the batch environment, the user ID of the batch job is used.

**Fast Path DEDBs**

In an online environment, the control region RACF user ID is used. (No batch environment exists.)

**Additional protection**

You can also implement database security with the DATABASE, FIELD, and SEGMENT classes in RACF.

**Related reading:** For more information on these resource classes, see “Preparing a RACF security plan” on page 332.

---

## Encryption: an alternative to access control

When preventing access to the data is difficult or impractical, encryption can protect data that is in files or data that is being communicated in a network. IMS offers some file encryption capability (through IMS Segment Edit/Compression exit routines, for example) but no communication encryption capability.

The following list provides additional information about encryption:

### Using cryptographic support

The Programmed Cryptographic Facility, program number 5740-XY5, provides file and communications encryption under z/OS. File encryption of the physical hierarchical database keeps unauthorized individuals from looking at the data when the physical disk pack containing the database is removed from its usual area. File encryption support extends to VSAM physical databases. Communications encryption supports ACF/VTAM supported terminals.

### Using the segment edit/compression exit routine (not DCCTL)

You can use this routine to provide data encryption. By including the IBM Programmed Cryptographic Facility within your exit routine, you can reduce your programming effort. The facility is executed by assembler macro calls. Segments are encrypted before being placed in the database buffer pool. The SEGM control statement in the IMS DBDGEN includes a keyword to specify the name of this exit routine.

### Using the ICSF/CCA interface

You can use ICSF/CCA APIs in the IMS DB Segment Edit/Compression exit. IMS supports the Programmed Cryptographic Facility (PCF) interface transparently through the ICSF/CCA interface. Programs that are written to the PCF interface run, without modification, through the ICSF/CCA interface. If you want your PCF programs to use the ICSF/CCA APIs, however, you must modify those PCF programs.

The ICSF/CCA interface has two PCF compatibility modes.

- ICSF mode COMPAT(YES) means that programs written to the Programmed Cryptographic Facility interface run without change, as well as calls made directly to the ICSF/CCA API. There are some limitations for dynamic master key change in this mode.
- ICSF mode COMPAT(NO) means only programs coded to the CCA API run.

---

## Security considerations for a DBCTL environment

This topic contains information about establishing security for an IMS DBCTL environment.

### DBCTL resources that can be protected

Before you decide what security facilities to use in designing a secure IMS system, you should know which resources within the system need protection. In other words, you should decide what to protect before you decide how to protect it.

The following resources can be protected in a DBCTL environment:

#### Control region

The IMS system control region that enables online application programs to process the database through terminals.

**System data set**

A collection of data that is fundamental to the operation of the IMS online system. An example is the IMS.PROCLIB data set.

**Dependent region**

An area of storage in the IMS online system in which batch or online application programs are executed. In a DBCTL environment, the dependent region can be a BMP region, a JMP region, or a CCTL region.

**PSB** Program specification block. The control block that describes a group of hierarchic databases and logical message destinations used by an online application program.

**BMP application program**

A program in a DBCTL environment that performs work for a user. Batch message programs are activated by the dependent region controller after the region is started by JCL.

**Database**

A collection of data that is fundamental to the user's activity. Using a Program Communications Block (PCB), a program has a logical view of the database, as described by the IMS physical database design.

## DBCTL security choices made during system definition

You make IMS security choices with initialization EXEC parameters or in two system definition macros: SECURITY and MSGEN. You can use the Resource Access Control Facility (RACF) to implement security decisions.

Use the specifications in the macro or initialization EXEC parameters to choose the type of security to be active in online execution. Use the Resource Access Control Facility (RACF) to name the resources.

If you use the SECURITY macro, you do not need to use the MSGEN macro to define security.

## Security facilities and security types for DBCTL resources

The security facility that you use to protect a particular DBCTL resource depends on the type of security that you want and the available security facilities for that type of security. Security facilities for DBCTL resources include z/OS, RACF, and PSBGEN.

The following table summarizes the types of security and facilities that are available for DBCTL resources.

*Table 27. DBCTL resources, security types, and security facilities.*

| Resources       | Type of security                                        | Security facility          |
|-----------------|---------------------------------------------------------|----------------------------|
| System data set | OS password protection                                  | z/OS                       |
|                 | Data set protection (VSAM)                              | RACF                       |
| Database        | Segment sensitivity                                     | PSBGEN                     |
|                 | Field sensitivity                                       | PSBGEN                     |
|                 | Password security (for the /LOCK and /UNLOCK commands). | RACF See note <sup>1</sup> |
| PSB             | resource access security                                | RACF                       |
|                 | APSB security                                           | RACF                       |

Table 27. DBCTL resources, security types, and security facilities (continued).

| Resources               | Type of security                                              | Security facility |
|-------------------------|---------------------------------------------------------------|-------------------|
| BMP application program | Password security (for the /IAM, /LOCK, and /UNLOCK commands) | RACF              |
|                         | Extended resource protection (using APPL keyword)             | RACF              |
| Control region          | Extended resource protection (using APPL resource class)      | RACF              |
| Dependent region        | RAS security                                                  | RACF              |
|                         | APSB security                                                 | RACF              |
|                         | resource access security                                      | RACF              |

**Note:**

1. RACF authority for the /LOCK PGM and /LOCK DB commands is not checked in a DBCTL-only system (SYSTEM parameter in the IMSCtrl macro during the system definition process contains DBCTL or IMS is started with the DBC procedure).

## Design considerations for DBCTL security

This topic explains how the various choices of IMS security can be used. When you are deciding on each part of your security design, consider the physical actions that an end user must take to obtain access to the system. You will probably use more than one type of security checking.

This topic assumes:

- A user identification as a control point
- The master terminal as a control point
- The use of RACF protection
- The use of a region as a control point

### Using password protection with command keywords

To provide verification before a command is accepted, you can require an accompanying password. The password is entered within parentheses immediately following the command verb.

### Limiting access from a dependent BMP, JBP, or CCTL region

Dependent BMP, JBP, and CCTL regions are resources you should protect. You can protect them by preventing the start of an unauthorized dependent region by start of task JCL and by preventing the use of unauthorized resources in a dependent region.

### Securing DBCTL dependent regions using RAS

RAS restricts the use of resources based on the user IDs of dependent regions. RAS uses RACF for security enforcement.

RAS allows an application program to access a PSB if the user ID of the dependent region in which the application program is running is authorized for that PSB. The

authority of a user ID to use a PSB is defined in either the IIMS or JIMS RACF security class. RAS does not restrict the application programs that can be scheduled in a dependent region.

The IIMS and JIMS classes are predefined by RACF; however, if they have not been included with your release of RACF, you can use the RACF ICHERCDE macro to define them as new classes.

You can specify RAS by specifying TYPE=RAS on the SECURITY system definition macro or ISIS=R execution parameter.

**Related reading:**

- For more information about the TYPE= keyword of the SECURITY macro and the ISIS= EXEC parameter, see *IMS Version 11 System Definition*.
- For more information about working with RACF security, see:
  - “Preparing a RACF security plan” on page 332
  - *z/OS Security Server RACF Security Administrator’s Guide*
- For more information about the ICHERCDE macro, see *z/OS Security Server RACF Macros and Interfaces*.
- For more information about updating the RACF resource class descriptor table, see *z/OS Security Server RACF System Programmer’s Guide*.

**Security considerations for Fast Path application programs**

When designing security protection for Fast Path application programs, or for DL/I programs that access Fast Path databases, consider that the processing in a dependent region can be protected by RAS security.

**Security for ODBA application programs**

You can secure any PSB specified on an APSB call from an ODBA application program using the z/OS System Authorization Facility (SAF) and an external security product, such as RACF.

For more information about security for ODBA application programs, see “Security for ODBA application programs” on page 322.

## **Activating IMS DBCTL security**

This topic gives guidance on the steps you take to activate your IMS security design using RACF and program exit routines.

Depending upon the security facilities you choose to use, you must perform the appropriate tasks.

- Define the SECURITY macro or initialization EXEC parameters for DBCTL.

The purposes of the SECURITY macro keywords that relate to DBCTL security are given here. Compare this discussion with the similar discussion in “Defining the SECURITY macro for DB/DC and DCCTL environments” on page 331.

The following SECURITY macro keywords control system security options:

**RCLASS**

The RCLASS keyword specifies an identifier to be used to identify the IMS system as a resource class to RACF for AGN verification. Resource class assignment is described in Table 28 on page 345.

**TYPE** Specifies whether application resource access authorization, transaction authorization, or signon user ID verification is to be performed. The



TYPE= keyword to specify which security facility, or combination of facilities, is to be included in your system.

**Related reading:** For more information on the SECURITY macro and its keyword parameters, see *IMS Version 11 System Definition*.

The following parameters specify various options related to implementing RAS security for the resources used by dependent regions:

#### **RASRACF**

For RAS security of dependent region resources using RACF alone. PSBs require an authorized dependent region's user ID before RAS grants access to them.

#### **RASEXIT**

For RAS security of dependent region resources using a user exit routine alone. PSBs require an authorized dependent region's user ID before the exit routine grants access to them.

**RAS** For RAS security of dependent region resources using RACF and a user exit routine. PSBs require an authorized dependent region's user ID before RACF and the exit routine grant access to them.

#### **NORAS**

RAS security of dependent region resources is not used.

- Prepare a RACF security plan in DBCTL.

To implement a RACF security plan:

1. Prepare a list of all the of IMS online resources to be protected, arranging them in groups to give an overview of the total resources covered.
2. Select the security facilities that protect the resource groups.
3. Code the SECURITY and MSGEN macros or initialization EXEC parameters.
4. Describe the resource class profiles to RACF.
5. Add users, groups, and data sets to RACF.
6. Modify JCL procedures in IMS.PROCLIB.

RACF resource classes are used by the IMS security function. The PSB class holds profiles for PSB security. RACF provides predefined resource classes or you can define your own. The names of resource classes to be used are specified with the RCLASS= parameter. Table 28 shows resource class assignments for DBCTL.

*Table 28. Resource class assignments for DBCTL*

| Resource class           | Resource class naming convention |                   |
|--------------------------|----------------------------------|-------------------|
|                          | RACF-defined name                | User-defined name |
| PSB resource class       | IIMS                             | Ixxxxxxx          |
| PSB group resource class | JIMS                             | Jxxxxxxx          |

The RACF resource classes are defined in RACF's resource class descriptor table (CDT). Initially, the IIMS and JIMS resource classes are predefined in the CDT. To add a resource class or to define resource classes with user-defined names, you must use the RACF resource class macro ICHERCDE to create an installation-defined CDT.

**Related reading:** For more information about the SECURITY macro and the RCLASS parameter, see *IMS Version 11 System Definition*.



## Controlling system startup for DBCTL security

The values generated for the DBC procedure all specify no security. You must reset them to enable security.

The ISIS execution parameter for the control region includes a way to control the kind of security checking that is done during the current execution. The parameter determines what flexibility the DBCTL operator must override the choice of security checking. You must coordinate the setting of the ISIS= keyword parameter with both overall security design and operational procedures.

The ISIS keyword parameter allows you to choose the type of dependent region security you want (RAS) and which security facilities, RACF or exit routines, will implement that security. You specify your choice using the ISIS= parameter. Your parameter options for ISIS= are as follows:

### **ISIS=0|N**

Disables RAS security.

### **ISIS=1**

Specifies RAS security using RACF. Specifying ISIS=1 is the same as specifying ISIS=N.

### **ISIS=2**

Specifies RAS security using an exit routine. Specifying ISIS=1 is the same as specifying ISIS=N.

### **ISIS=R**

Specifies RAS using RACF.

### **ISIS=C**

Specifies RAS using an exit routine.

### **ISIS=A**

Specifies RAS using both RACF and an exit routine.

The following table shows the SECURITY macro parameters that the ISIS= keyword overrides or replaces.

*Table 29. The ISIS= keyword and the SECURITY macro parameters it overrides*

| JCL parameter and Value | Security macro parameter it overrides     |
|-------------------------|-------------------------------------------|
| ISIS=R                  | TYPE=NORAS, TYPE=RASEXIT<br>TYPE=RAS      |
| ISIS=C                  | TYPE=NORAS, TYPE=RASRACE,<br>TYPE=RAS     |
| ISIS=A                  | TYPE=NORAS, TYPE=RASRACE,<br>TYPE=RASEXIT |
| ISIS=N                  | TYPE=RASRACE, TYPE=RASEXIT,<br>TYPE=RAS   |

If RACF is to be used for a CCTL, the CCTL JOB statement must include the USER specification. The CCTL itself can provide security checking (RACF or its own) independent of the DBCTL security checking that is described in this section.

## Implementing DBCTL security changes online

If you use online change for system definition, you must update the security definitions for RACF.

As a result of online changes to the system definition, you might need to take the following steps to keep your security specifications current:

- Update passwords.
- Add security provisions to the online system for password security.
- Refresh the RACF database.

To change RACF security definitions online, update the RACF database. Then issue the RACF command SETROPTS RACLIST(*classname*) REFRESH.

## Controlling DBCTL security violations

Security violations are handled according to the installation's security administration guidelines.

IMS records the following security violation attempts on the IMS system log:

- Password omitted when one is required
- Password incorrect for authorization
- Misspelled password

Any of these errors causes IMS to log the violation. The IMS system log provides an audit trail for investigation of possible security problems. The IMS system log security violation is identified as a X'10' log record type. You can use the File Select and Formatting Print utility to print the log.

Another method for recording security violations is available when RACF is installed. Each resource access violation creates a RACF type 80 record.

### Related reading:

- For more information on the File Select and Formatting Print utility, see *IMS Version 11 System Utilities*.
- For more information on utilities to format and print RACF records, see *z/OS Security Server RACF Auditor's Guide*.



---

## Chapter 19. Controlling IMS

Controlling IMS consists of many tasks.

These tasks are discussed in the following topics.

**Related reading:** For more detailed information, see *IMS Version 11 Operations and Automation*.

---

### Monitoring the system

Monitor the status of the system on a regular schedule to gather problem determination and performance information.

For example, to determine if you should start an extra message region, you might monitor the status of the queues during peak load.

**Related reading:** For more information on monitoring, see Chapter 23, “Collecting and interpreting IMS monitoring data,” on page 369.

---

### Processing IMS system log information for analysis

The system log data sets are a basic source for statistics about the processing performed by the online system. Individual log record types contain data that can be analyzed in many ways.

For example, you can select and format all activity pertaining to a specified user ID or about IMS pools.

### Using IMS system log utilities

IMS provides several utilities to assist with extracting log records from the system log. These utilities also assist with reducing and merging data that is spread across several log data sets.

The sections that follow describe several of these utilities:

- **File Select and Formatting Print utility**

You can use the File Select and Formatting Print utility (DFSERA10) if you want to examine message segments or database change activity in detail. This utility prints the contents of log records contained in the OLDS, SLDS, or the CQS log stream. Each log record is presented as one or more segments of 32 bytes. The printed output gives each segment in both character and hexadecimal formats.

- **Fast Path Log Analysis utility**

Use the Fast Path Log Analysis utility (DBFULTA0) to prepare statistical reports for Fast Path, based on data recorded on the IMS system log.

These reports are useful for system installation, tuning, and troubleshooting. The Fast Path Log Analysis utility is not related to the IMS Monitor or the Log Transaction Analysis utility. For more information about the IMS Monitor, see “IMS Monitor” on page 378.

- **Log Transaction Analysis utility**

In an IMS DB/DC or DCCTL environment, you can use the Log Transaction Analysis utility (DFSILTA0) to collect information about individual transactions, based on records in the system log. Many events are tabulated in the Log Analysis report that is produced by this utility, including total response time, time on the input queue, processing time, and time on the output queue.

- **Statistical Analysis utility**

In an IMS DB/DC or DCCTL environment, you can produce several summary reports using the IMS Statistical Analysis utility (DFSISTS0). You can use these reports to obtain actual transaction loads and response times for the system. The statistics produced are dependent on the input system log data sets.

- **Knowledge-Based Log Analysis**

Knowledge-Based Log Analysis (KBLA) is a collection of IMS utilities that select, format, and analyze log records. KBLA also provides an ISPF interface to create and run the jobs for various log-related utilities, and to access other ISPF applications. KBLA uses an ISPF panel-driven user interface to simplify JCL job creation and to prevent JCL errors. KBLA generates the JCL and control statements necessary to run the supported utilities. This JCL preparation allows you to focus on the output of the utility used rather than on how to code JCL to extract information.

**Related concepts:**

“Analyzing log records using KBLA utilities” on page 833

**Related reference:**

 Knowledge-based log analysis utilities (System Utilities)

## IMS database productivity tools

IBM offers a number of IMS database productivity tools. IMS Tools is a set of database performance enhancements for your IMS environment.

These tools can help you automate and speed up your IMS utility operations. They can also assist you in analyzing, managing, recovering, and repairing your IMS databases.

**Related reading:** To learn more about these tools, visit the DB2 and IMS Tool web site at <http://www.ibm.com/software/data/db2imstools/>.

---

## Recovering the system

Ensure that resource definitions have been exported to a resource definition data set (RDDS) if you use DRD and if you perform a cold start from an RDDS.

You must ensure that any resource changes made dynamically are recovered across a cold start.

While IMS is running, an IMS system programmer or operator might need to complete tasks to recover the system.

**Related concepts:**

 Overview of dynamic resource definition (System Definition)

**Related tasks:**

 Executing recovery-related functions (Operations and Automation)

---

## Modifying and controlling system resources

You establish the initial settings of IMS resources as part of the IMS system definition process. The MTO, and other operators authorized to do so, can change various system resources using IMS commands.

You can modify and control the operating state of the following IMS system resources using type-1 commands:

- Dependent regions
- Telecommunication lines
- Terminals
- Transactions
- Databases
- ISC users (subpools)
- ETO users
- MSC resources
- Security options
- Conversations
- Subsystems

You can use many IMS commands to perform similar control functions for different types of resources.

The resources are:

- Telecommunication line, physical terminal, or node resources
- Logical terminal resources
- Logical link resources
- Logical link path resources
- Transaction resources
- Transaction class
- Program resources
- Database resources
- Subsystem resources

The *IMS Version 11 Operations and Automation* show the relationship between these commands and resources, and provide answers to a series of specific questions. For example, after a command is issued, can a resource:

- Receive input?
- Send output?
- Perform output message queuing?

You can also modify and control IMS resources using type-2 commands in:

- Single-IMS IMSplexes that are configured with a minimal Common Service Layer that include an Operations Manager (OM) and a Structured Call Interface (SCI), but no Resource Manager (RM). This configuration is also known as an enhanced command environment.
- Multiple-IMS IMSplexes that are configured with a full CSL (OM, SCI, and RM).

**Related reading:**

- For details about the system definition processes (batch and dynamic) and details about configuring a single IMS with a minimal CSL, see *IMS Version 11 System Definition*.
- For details about the IMS commands, see *IMS Version 11 Commands, Volume 1: IMS Commands A-M* and *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

---

## Controlling data sharing

Controlling data sharing involves using DBRC and monitoring the data-sharing system.

### Controlling data sharing using DBRC

DBRC allows you to control access to data by IMS subsystems that participate in data sharing. Using DBRC, you can modify, initiate, and delete the current status indicators in the RECON data set to change the access intent of online IMS subsystems and the share level of registered databases.

Your data sharing environment depends on the status of the databases and subsystems as indicated in the RECON data set.

**Related reading:** For more information about how DBRC participates in data sharing, see Chapter 34, “Supporting data sharing,” on page 499.

You can modify the share level indicator using a form of one of the DBRC online change commands, /RMCHANGE.

### Monitoring the data-sharing system

To monitor data sharing, you obtain information on the status of the IRLM, IMS subsystems and databases, the RECON data sets, and, for data sharing in an IMSplex, coupling facility structures that are participating in data sharing.

To:

- Display the status of an IRLM on either your system or on another connected system, enter the following z/OS command:  
`MODIFY irImproc,STATUS,irlmx`  
**Related reading:** For a complete description of the commands for the IRLM, see *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*.
- Monitor components and resources in a data-sharing environment, use the same types of procedures as in a non-data-sharing environment. For more information about monitoring, see Chapter 23, “Collecting and interpreting IMS monitoring data,” on page 369.
- Monitor structures on a coupling facility, use following z/OS operator commands:
  - `DISPLAY XCF,STRUCTURE`: Used to display the status of structures defined in your active policy.
  - `DISPLAY XCF,STRUCTURE,STRNAME=`: Used to display detailed information about a specific structure. The structures that are important to IMS are the IRLM, OSAM, VSAM, shared MSGQ, and shared EMHQ structures.

These commands let you look at structures on a coupling facility to determine resource status and, for failures, gather information for problem determination.

**Related reading:** For detailed information about these commands, see *z/OS MVS System Commands*.

---

## Controlling log data set characteristics

From time to time, you need to tune and modify log data set characteristics.

For example, in the following circumstances you should tune and modify log data set characteristics:

- After monitoring
- After changing your requirements for system availability, integrity, or operator handling

For detailed information about controlling the OLDS, WADS, SLDS and RECON data sets, see *IMS Version 11 Operations and Automation*.

---

## Connecting and disconnecting subsystems

Before an IMS subsystem can access external subsystems (another program executing in an z/OS address space) such as DB2, you must connect the IMS subsystem to this other subsystem. IMS can connect to another subsystem only if that subsystem is identified in the subsystem member in IMS.PROCLIB.

Specify the subsystem member name in the SSM EXEC parameter. When specified to and accessed by IMS, the subsystem member name cannot be changed without stopping IMS.

**Related reading:** For information on the SSM EXEC parameter, see *IMS Version 11 System Definition*. For more information about connecting IMS to another subsystem, see *IMS Version 11 Operations and Automation*.





---

## Chapter 20. Starting or restarting IMS

Starting or restarting an IMS subsystem means initializing the IMS control region.

To restore the operating environment, the MTO must ensure the following are also initialized:

- All IRLMs
- All CQSs
- Any dependent regions (including utilities)
- All CSL address spaces (OM, RM, and SCI), if operating as an IMSplex using CSL
- Any connections to VTAM
- All communication lines and terminals

The term MTO used in this information refers to the IMS master terminal operator for the DB/DC and DCCTL environments, and to the operator for the DBCTL environment. All commands shown assume that the command-recognition character is the forward slash (/).

See *IMS Version 11 Operations and Automation* for additional information.



---

## Chapter 21. Shutting down IMS

The command that is used to shut down the control region also forces termination of data communications and the dependent regions if they have not already been terminated in an orderly way, and can also tell the CQS subsystem to shut down.

A common sequence for shutting down the entire online system is:

1. For an IMS DB/DC or DCCTL environment, stop data communications. For an IMS DBCTL environment, disconnect from the CCTL.
2. Stop dependent regions.
3. Stop the control region.
4. For an IMS DB/DC or DBCTL environment, stop the IRLM.
5. For a shared-queues environment, if CQS has not shut down, shut it down.
6. For an IMSplex system using CSL, shut down the CSL manager address spaces (OM, RM, and SCI).

### Related tasks:

 Shutting down IMS (Operations and Automation)



---

## Chapter 22. Testing the system

As a system administrator, you must be involved in two phases of testing. One phase occurs when application programs are verified prior to a cut over to production mode. The second phase involves evaluating changes or corrections to existing application programs to ensure that application function has not regressed and to validate the new or changed function.

When you have multiple application programs, you must ensure that modifications to one application program do not impact the service provided to any end users.

The various testing phases, with the corresponding administration tasks and related development activity, are listed in the following table.

*Table 30. Administration tasks related to testing phases*

| Test phase            | Administration task                          | Related development activity                         |
|-----------------------|----------------------------------------------|------------------------------------------------------|
| Unit test             | Identify bottlenecks                         | Test development code                                |
| Function test         | Plan operations procedures                   | Test explicit functions                              |
| Integration test      | Prepare test system definition               | Build the system                                     |
|                       | Plan test database                           | Build test database                                  |
| Component test        | Verify network operation                     | Validate major portions of application program logic |
| System test           | Check out operations and recovery procedures | Build fully executable system                        |
|                       | Coordinate use of test tools and monitoring  | Validate operational procedures                      |
|                       | Ensure network readiness                     | Test coexistence                                     |
| Performance testing   | Plan for simulation                          | Validate claims and set benchmarks                   |
| Stress testing        | Plan for peak loads and response criteria    | Test for high volume of traffic and processing       |
| Acceptance test       | Finalize operations procedures               | User liaison checkout on behalf of end users         |
| Maintenance testing   | Control libraries and online definition      | Application and IMS service checkout                 |
| Design change testing | Plan response across administration tasks    | Control verification of changes                      |
| Regression testing    | Plan monitoring                              | Verify that old function is not damaged              |

---

### The need for a test system

The development of a test system requires the participation of representatives from several areas.

Your primary concern at the final stage of implementation is to offer satisfactory service to the end users. Does the IMS system perform as expected? A secondary

concern, when changes occur in one or more online application programs, is to preserve the integrity and service for all end users. Will changes seriously impact the production environment?

Your installation might have a separate test organization. In this case, development personnel do not test the working code but might have to demonstrate it before handing over the application programs for independent testing.

The end user might be represented by a user-liaison group that knows the business needs. The liaison group also might evaluate the adequacy and accuracy of application programs.

Your solution to the need for a testing system can take several forms:

- A separate IMS test environment that is operational during major development activity
- A separate IMS test system used for ongoing verification of maintenance and application design changes
- A production system that allows controlled changes to be tested online, possibly using regions that are initialized for test purposes at a time that does not impact production processing

Part of your administration role is to ensure that a suitable online IMS test system is created and that the test procedures and any special database requirements are well documented. This is important if the test system is also to function as the maintenance system. However, your role can be more passive, in which you only participate in testing to verify operations procedures and to verify that all system definitions and preparation for production mode are in place.

After the application package has entered the production phase, the role of the developer or tester is usually taken over by program maintenance personnel. Administration needs to assess the impact of maintenance to the existing online IMS design. This task is described in more detail in Chapter 24, "Modifying the system design," on page 395.

## Setting up a test system

An administration task is to define the procedures for applying changes to the online IMS system. Part of this task is to have an acceptance procedure that is followed by all responsible participants.

The major considerations for establishing a separate test system are:

- The installation's policy for protecting the integrity of the database
- The presence of criteria for accepting an application as properly tested
- The need for a stable production environment
- The degree of support available to solve day-to-day problems
- The complexity of the applications
- The degree to which the current programs can perform correctly, or function with interim problem bypasses.

Some of the factors involved in the decision to use a test system as an ongoing function are:

- Will there be staged implementation of new application programs or add-on function?

- Can online service interruptions caused by the need to fix application problems encountered during testing be tolerated and, if so, to what degree?
- What procedures should be followed for accepting system changes or corrections?
- How is maintenance to the IMS system itself to be handled—will the test system be used?
- Will corrections to the application programs and IMS problem fixes be incorporated on demand or batched at agreed-upon intervals?

## Setting up a test database

A test database must contain enough data to adequately test the system. The data should exercise a major portion of an application program's logic to prevent regression of existing function and yet economize on the amount of data.

Usually, a test transaction stream executes against a known database status, and the results are compared to predicted results. The content of the transactions and implications of the database status needs to be analyzed for accuracy by a user-liaison, group as well as by groups responsible for the database design. As additional test cases are added, they must be validated and inspected for redundancy. The trade-off is between adding to an existing transaction or defining an additional transaction that might require extra database content. The test case stream must be adequately documented, so that you know what it does and what data it needs.

**Related reading:** For more information on creating test databases, see *IMS Version 11 Database Administration*.

## Testing operational procedures

You can use the test system to ensure the accuracy and usability of your operational procedures.

The following activities are suitable during the system test phase:

- You can test preliminary versions of the MTO procedures, the run book, and the incident report forms. The MTO can perform system startup, connect a subset of the terminals and nodes, and practice the restart actions.
- You can hold an informal audit of command use and discuss any misunderstandings to help you refine the content of the operator instructions.
- During online execution, you can arrange other events that require MTO intervention to take place, such as taking image copies, invoking the IMS Monitor and traces, and responding to application program abnormal termination.
- You should also test the recovery procedures. Have the operations staff carry out a database recovery and verify the correct use of system logging control and recovery utilities. You can build a test RECON data set for DBRC and follow through the GENJCL step to recover to a given checkpoint position.

If the test system is used for checkout of changes to the application package, any significant operational changes should also be exercised, possibly with MTO observers.

## Monitoring in IMS test environments

Your objectives for monitoring during a test phase are slightly different from those for production systems



Three activities you should plan for are:

- Detecting and correcting potential performance problems before entry into production mode.

Using realistic data, you can expose patterns of processing for the application programs, especially the DL/I call occurrences. Compare the Monitor results against the expected transaction profile. Such items as excessive I/O events or large I/O wait times can reveal a performance problem at an early stage.

- Testing the monitoring part of your operations procedures.

Become familiar with the way output is produced and the format and content of reports. Start to develop work sheets to summarize Monitor findings rather than write unstructured comments on the report output.

**Related reading:** See Chapter 23, “Collecting and interpreting IMS monitoring data,” on page 369 for information about the detailed report descriptions and considerations and tools for your monitoring strategy.

- Using your monitoring tools for base profiles of new program processing.

Discuss your findings with development personnel and with performance specialists. Try to obtain early warning of performance problems. If you are integrating a new application into an existing online IMS system, try to assess beforehand the impact of the added application workload.

During the testing stage you can also perform stand-alone or calibration runs to establish base profiles for the critical transactions. The Call Summary report from the IMS Monitor is useful for this purpose. You can also compare the results for the tested transactions to similar transactions already in the production system.

## Monitoring in a DB/DC environment

In a DB/DC environment, you can obtain information about the precise sequence of calls issued by application programs using the report capabilities of IMSASAP II.

This tool requires using the online IMS Monitor to produce the monitor trace records. Its use as a monitoring tool and installation prerequisites are described in Chapter 23, “Collecting and interpreting IMS monitoring data,” on page 369.

---

## Ensuring network readiness

When you prepare an application package for production mode, you must be aware of the status of all terminals or connected devices planned for availability with the online IMS system.

- Develop a detailed implementation plan for each device and control unit. Some of the items in this document might be:
  - Exact type of device, model, and operational characteristics
  - The configuration for a set of components
  - Physical location and person responsible for the device installation
  - The names by which the online IMS system knows the device: LTERM, line number, unit address, and node names
  - The correct VTAM MODEENT macro's PSERVIC parameters for: LUNAME, LUTYPE, TS Profile, 32xx model, 3270 screen size, and NTO device type
  - The logon, user, and MSC descriptors for terminals and users defined with ETO
  - The appropriate exit routines

- If an output device, the source of paper supplies and arrangements for distribution of output
- If already in operation outside of IMS, the restrictions on the use of the device for test purposes
- System definition requirements for the device, and the planned timing of the generation
- If the device is programmable, what programs are necessary, locally and in the host, for IMS execution
- Use the terminal profiles built up as part of the preparation for system and network definition. You need this information because of the potentially large number of devices that can operate in the online IMS system. Also, you must converse in terms of the hardware with installation personnel and data communications specialists whose responsibilities usually lie beyond IMS device support.
- Contact the person responsible for the system definition regarding the status of statically defined VTAM terminals. The definition events are usually scheduled as part of a master plan for implementing the application package.

## **Network testing in DB/DC and DCCTL environments**

As part of your preparation for production mode, you should arrange for online testing of each piece of the network during the system test. Try to observe the online testing for the different terminal types that are to be active in the IMS system.

Include observations of those transactions considered critical. The results of this exercise are useful when you write operations procedures for remote terminals and instructions for the master terminal operator. Emphasizing the physical actions and sequence of events helps clarify these procedures and make them more complete for the end user.

If you do not have access to an early draft of the actual procedure for the terminal operation, it is best to use a prepared script for a session. The text should include a scenario depicting a test sequence of commands as an online session:

1. Prepare to transmit the SNA terminal subsystem program, then transmit from the host (if applicable).
2. Start the ACF/VTAM or host subsystem application program, and enter the appropriate commands to start IMS and establish communication with VTAM.
3. Start the SNA terminal subsystem program at the processing unit location (if applicable) and make the intelligent terminal ready for communication with IMS.
4. Exchange messages with the host, submit an IMS transaction from the terminal, and validate the result.

## **Testing in a DBCTL environment**

If you are testing a DBCTL environment with a CCTL, be aware that the CCTL, not IMS, controls the network, terminals, and transactions.

## **Testing in a DCCTL environment**

You can test any DCCTL function that does not require access to a database.

For example, you can schedule transactions or perform component testing if no database calls are made. If your test system already has a GSAM or an external subsystem (for example, DB2 for z/OS) installed, you can verify connections and applications requiring those systems.

---

## IMS testing aids

The following topics provide information about various testing aids.

**Related reading:** For more information on suggested procedures and utilities used during testing of application program code, see “Testing an IMS application program” in *IMS Version 11 Application Programming*.

### Simulating online execution with Batch Terminal Simulator in DB/DC and DCCTL environments

Often an online application program is developed to the testing stage, but a suitable online IMS system is not available. At this stage you can use the Batch Terminal Simulator licensed program, program number 5655-J57.

The Batch Terminal Simulator (BTS) program can simulate the operation of message processing regions before the online IMS system is operational as a test system. This program runs as a batch IMS system using one or more applications. BTS input is transaction data, and the application programs are invoked. Database calls are executed against test databases; the DL/I calls for data communications are simulated.

With BTS, you can do the following:

- Print terminal input, output, and an optional trace of related database activity. You can request a summary of DL/I calls, by type, against each PCB.
- Simulate conversational transactions and program-to-program switches.
- Simulate the message queuing and application program scheduling functions of the online IMS system. transaction input data to the input message structure for the application program and, similarly, map the output messages to a printed layout. This is particularly useful for IMS 3270 input and output formats.
- Use the debugging and trace facilities (during the test phase) for both batch and telecommunication applications.

You do not need to modify the IMS control programs, control blocks, libraries, or the application programs.

**Related reading:** For more information on BTS, see *IMS Batch Terminal Simulator for z/OS User's Guide*.

### Online testing of MFS formats in DB/DC and DCCTL environments

If you need to test an application program change online in either a test or a production system and the test transaction input uses an MFS-supported screen format, you can use the MFSTEST mode.

IMS allows individual MFS-supported terminals to enter MFSTEST mode. In this mode, you can use temporary message formats from an alternative MFS format

library, rather than changing the message format blocks in the production format library. A format block name can be identical to the name of a message format already in the production library.

An MFS-supported terminal is placed in test mode by entering a `/TEST` command with the MFS parameter. Next, the remote terminal operator enters a `/FORMAT` command for transaction processing, causing message formats to be selected from a library of test formats (IMS.TFORMAT). If these message formats are not found in this library, they are selected from the active message format library (IMS.FORMATA/B). By entering an `/END` command, the terminal is removed from test mode.

When testing MFS formats, you can use the `/TRACE SET ON TRAP` command to trap and analyze MFS errors. This command helps you analyze errors that result from incorrect manipulation of the MFS control blocks.

### **System definition requirements for MFSTEST mode**

To use MFSTEST formats, specify the appropriate parameter on either the `IMSGEN` or `COMM` macro.

**Related reading:** For more information on the `IMSGEN` and `COMM` macros, see *IMS Version 11 System Definition*.

### **Online execution requirements for MFSTEST mode**

When you want to perform testing of MFS formats during execution of the online IMS production system, the `IMS.TFORMAT` library must be defined and must contain the appropriate test blocks.

These can be named the same as existing production formats. The `IMSTFMTA/B` DD statements must be present in the control region JCL, with the `IMS.TFORMAT` library as the first data set, followed by the corresponding `IMS.FORMATA/B` library concatenated to it.

The `CIOP` is used to hold all format blocks for terminals operating in MFSTEST mode. Storage space in `CIOP` is dynamically allocated according to your system demands.

Executing in MFSTEST mode does result in a performance impact, because both transmission and message formatting are being stored in the same buffer.

## **Using dynamic resource definition and online change for testing**

Using dynamic resource definition and online change capabilities, you can plan to perform selective testing during execution of an IMS online test system. Your testing can be accomplished without generating a modified test system merely to accommodate some temporary changes.

You can perform testing that supports:

- The addition of new applications
  - New transactions, programs, and MFS formats
  - Additional databases
- Modifications to existing applications
  - Modified database definitions
  - Replaced PSBs and programs

- Altered MFS formats
- Necessary deletions for the above modifications

You cannot use this method to test the use of terminals or devices not already defined to the IMS system.

Another factor might be the security arrangements. Planning for change must ensure that appropriate authorization exists to use the transactions or existing terminals for the duration of the test.

You might plan for a group of application changes to make up a test package to be executed during IMS online operation in a test system, or even during a production cycle. Online change is best used to migrate pretested changes without the need for a full restart interruption. Your planning should include the following considerations:

- Online data resources must be adequately protected. Carefully assess the risk that an application program being tested could damage production data.
- Added databases must be included in the JCL for the IMS control region; a solution for z/OS systems is to use dynamic allocation.
- Although changes to database structures can be reflected in a new ACB build, this might not always be practical if the changes require a reorganization.
- You must coordinate the update of IMS.PGMLIB with the status of the online processing. A replaced program might cause problems, unless the end user is aware of the activity. The library update might have to be carried out in the interval just before the /MODIFY COMMIT command.
- You must assess the potential performance impact to tuned systems, although a monitored test execution can provide valuable information for performance planning.

The entry of the /MODIFY PREPARE command begins the cut-over operation. When activity has stabilized for the resources affected by the content of the online change, the /MODIFY COMMIT command completes the cut-over. Testing can now begin. When you complete the planned testing, repeat the /MODIFY PREPARE, /MODIFY COMMIT command sequence to reinstate the old unmodified system data sets that were inactive during the testing. It is the responsibility of those performing the test to nullify the effects of the testing and back out any inappropriate changes made to the databases.

## **Program testing using SYSIN/SYSOUT in DB/DC and DCCTL environments**

One way of testing a message processing program is to provide an input data stream containing messages.

Specify a line group as a READER, then assign the input SYSIN to a local card reader. Messages are passed to the application program by GU calls to the message queue. No editing or logging of position occurs while messages are being processed. The accuracy of the end-of-message or segmentation is determined by the input stream.

Similarly, for output, you can assign a line group as a printer, punch, tape, or DASD device using the UNITYPE keyword for the LINEGRP macro. You can then assign the output LTERM for the appropriate device characteristics, although the actual device does not have to be allocated.

The LINEGRP macro parameters produce a DD statement for the control region, and the appropriate buffer size is defined for the output records. For printer output, a translation to the 48-character set, lowercase to uppercase, occurs; all other codes are converted to periods.

## **Network testing using Teleprocessing Network Simulator in DB/DC and DCCTL environments**

Teleprocessing Network Simulator (TPNS), program number 5662-262, is a terminal and network simulation tool used for determining system performance and response times, evaluating teleprocessing network design, functional testing, and automating regression test procedures.

## **Performance and stress testing using the Queue Control Facility**

Queue Control Facility (QCF) can be used in DB/DC and DCCTL environments to select messages from actual workload cases saved in archived log data sets. You can re-queue the messages at selected rates to perform stress and performance testing.

**Related reading:** For more information on the Queue Control Facility, see *IMS Queue Control Facility for z/OS User's Guide*.



---

## Chapter 23. Collecting and interpreting IMS monitoring data

Monitoring is the collection and interpretation of IMS data. This data helps you understand the daily needs of the system and provides an insight into areas that might need changes.

Monitoring should be an ongoing task because:

- Monitoring helps you establish base profiles, workload statistics, and data for capacity planning and prediction.
- Monitoring gives early warning and comparative data to help you prevent performance problems.
- Monitoring validates tuning you have done in response to a performance problem and ascertains the effectiveness of that tuning.

For multi-system networks, plan to obtain both statistical and performance data for IMS online systems that are part of the network. You can use the same monitoring tools that are used for generating performance data for single IMS systems:

- The IMS Monitor can be executed concurrently in several systems. You obtain IMS Monitor reports for each individual IMS system and coordinate your processing analysis.
- For DB/DC and DCCTL systems:
  - The IMS Statistical Analysis utility (DFSISTS0) produces summaries of transaction traffic for each individual system. Again, you combine the statistics for a composite picture.
  - The IMS Log Transaction Analysis utility (DFSILTA0) enables you to trace transactions across multiple systems and examine the traffic using the various active physical links.

Overall, conclusions from continuous monitoring and historical records provide a good starting place from which to answer end-user complaints and to provide initial directions for tuning projects.

---

### Establishing monitoring procedures

Several types of monitoring strategies are available.

You can:

- Summarize actual workload for the entire online execution. This can include both continuous and periodic tracking. You can track total workload or selected representative transactions.
- Take sample snapshots at peak loads and under normal conditions. It is always useful to monitor the peak periods for two reasons:
  - Bottlenecks and response time problems are more pronounced at peak volumes.
  - The current peak load is a good indicator of what the future average will be like.
- Monitor critical transactions or programs that have documented performance criteria.



- Use the z/OS Workload Manager to help manage workload distribution, balance workloads, and distribute resources.

Plan your monitoring procedures in advance. A procedure should explain the tools to be used, the analysis techniques to be used, the operational extent of those activities, and how often they are to be performed.

Regardless of which strategy you use, you need to:

- Develop performance criteria
- Develop a master plan for monitoring, data gathering, and analysis

## Establishing performance objectives

Inherent in the design of your online IMS system are your performance objectives. Establishing performance objectives is a major task requiring data gathering for the IMS workload as a whole.

After defining the workload and estimating the resources required, you must reconcile the desired response with the response you consider to be attainable. Monitor the performance of the system to determine if these objectives are being met. Base performance objectives on:

- Desired, acceptable, and maximum response time
- Average and maximum resource demands (or workload) per transaction
- Predicted and actual transaction volumes

Establishing your performance objectives is an iterative process involving the following activities:

1. Defining user-oriented performance objectives and priorities.

These derive from the way an end user perceives the service provided by the system. For IMS these objectives state expectations of response time as seen by the end user at the terminal.

When response time objectives are established, they should not only reflect the time-in-system (the elapsed time from entry of a last input message segment to the first response segment), but also the expected amount of IMS and application program processing. You should consider whether to define your criteria in terms of the average, the 90th percentile, or even worst-case response time. Your choice depends on your installation's audit controls and the nature of the specific transactions.

2. Determining how performance against these objectives is measured and reported to users.

This includes identification of systematic differences between the measured data and what the user sees. You should investigate the differences between internal (as seen by IMS) and external (as seen by the end user) measures of response time. To do this, you can use the following tools:

- The Transaction Response report produced by the IMS Statistics Analysis utility, which gives the following data on internal response time by transaction type:
  - Longest and shortest response
  - 25th, 50th, 75th, and 90th percentiles of the response time distribution
- Installation-written programs, which can analyze the output of the IMS Log Analysis utility. Other programs can also provide the same type of information, and are usually tailored to satisfy the installation's requirements.

3. Understanding and documenting the current workload.

This requires breaking down the total work into categories and, for each category, developing a workload profile (generally estimates) that include:

- Definition of the transaction category (for example, transaction type or group of transactions). The two characteristics of the category are:
  - The IMS workload, which is generally documented by transaction profile. In a well-designed IMS online system, most transactions perform a single function and have an identifiable workload profile. Later in the process, transaction types with common profiles can be amalgamated for convenience.
  - The transaction volume. In situations where the workload to be documented is already operational, a summary of transaction volumes can be obtained from the IMS Statistical Analysis utility. (The Application Accounting Report gives transaction counts within program name.) In other cases, volumes are estimated.
- Relative priority of category, including periods during which priority changes.
- Resource requirements of the work:
  - Physical resources managed by the operating system (real storage, DASD I/O, terminal I/O)
  - Logical resources managed by the subsystem, such as control blocks, latches, buffers, and number of regions

To obtain a base profile of transaction resource demands, you can start IMS on a dedicated machine and execute a few transactions to accomplish initialization and buffer pool usage. Then start the IMS Monitor and measure a sample of transaction execution.

You can use the base transaction profile to examine the transaction workload to see if it can be reduced. Such design changes result in the greatest impact, occurring before system-wide contention. You can also compare the base profile to the transaction profile in the production environment.

4. Translating the resource requirements and volume information obtained into system-oriented objectives for each work category

This includes statements about the transaction rates to be supported (including any peak periods) and the internal response time profiles to be achieved.

5. Confirming that the system-oriented objectives are reasonable.

After initializing the system and monitoring its operation, you need to find out if the objectives are reasonable (given the hardware available), based upon the measurements of the workload. If the measurements differ greatly from the estimates used, you must revise the workload documentation and system-oriented objectives accordingly, or tune the system.

6. Establishing performance objectives.

Establishing performance objectives is also a necessary prerequisite to using the z/OS Workload Manager. Much of the information gathered when you establish performance objectives can be used as input when planning for workload management.

## Planning for workload management

z/OS provides a workload management function to help you manage workload distribution, balance workloads, and distribute resources to competing workloads.

z/OS provides this support automatically after you specify in the panel-driven application that the z/OS Workload Manager (WLM) provides how you want your workloads processed.

# Workload Manager and IMS

With the z/OS WLM, you define to z/OS the performance goals for transactions and the relative importance of transactions and address spaces. Using your definitions, WLM then decides how the resources that are controlled by z/OS should be allocated.

## Performance goals and business importance for transactions:

A performance goal can be the average response time desired for a transaction, or it can be that a certain percentage of the transactions complete within the response time period.

The business importance is a priority level representing how critical a type of transaction is to your installation. 1 is the highest priority level.

After a transaction is scheduled by IMS, the z/OS WLM uses the performance goals that you define for each transaction to decide how much resources, such as processor cycles and storage, it should allocate to meet your goals. When contention for system resources occurs, the z/OS WLM uses the business importance assigned to the transaction to help decide which transactions get priority for the resources that are under the control of the z/OS WLM. The performance goals and business importance are defined in a service definition.

**Related reading:** For more information about service definitions, see “Establishing the service definition” on page 373.

## Assigning relative importance to IMS address spaces:

After setting performance objectives for IMS transactions, specify in WLM service classes the importance of each IMS address space relative to all other IMS address spaces. Using your specifications, WLM determines the processor dispatching priority for all address spaces in the z/OS system.

In the following figure, the WLM ISPF-driven panel for defining a WLM service class assigns relative importance of 2 to the IMS control region address space.

MODIFY A SERVICE CLASS

ROW 1 TO 00000100

COMMAND ==> 00000200

00000300

SERVICE CLASS NAME . . . . . : CNTL 00000400

DESCRIPTION . . . . . IMS CTL RGN SERVICE CLASS 00000500

WORKLOAD NAME . . . . . IMS (NAME OR ?) 00000600

BASE RESOURCE GROUP . . . . . (NAME OR ?) 00000700

CPU CRITICAL . . . . . YES (YES OR NO) 00000800

00000900

SPECIFY BASE GOAL INFORMATION. ACTION CODES: I=INSERT NEW PERIOD, 00001000

E=EDIT PERIOD, D=DELETE PERIOD. 00001100

00001200

---PERIOD--- -----GOAL----- 00001300

ACTION # DURATION IMP. DESCRIPTION 00001400

— 1 2 EXECUTION VELOCITY OF 70 00001500

00001600

00001700

Figure 42. Assigning importance using the WLM ISPF panel Modify A Service Class

The following table suggests how you might assign WLM importance values to achieve a recommended relative prioritization of some of the IMS address spaces. The N in the first column of the table represents a value from 1 to 5. 1 is the highest priority.

*Table 31. Assigning WLM importance to IMS address spaces*

| WLM importance | IMS address space                                 |
|----------------|---------------------------------------------------|
| N              | IRLM, DBRC                                        |
| N + 1          | IMS control region, CQS, OM, RM, SCI, IMS Connect |
| N + 2          | DL/I separate address space                       |
| N + 3          | Dependent regions                                 |

By prioritizing IMS address spaces as shown in the table, you minimize the risk of critical functions having to wait for dependent regions to finish processing before they receive the processor resources they need. Such critical functions include locking, DB authorization, IMSplex-wide command processing, messaging, and DB open and close activities.

**Related reading:** For additional information on dispatching priorities, see “Dispatching priorities” on page 411.

## Establishing the service definition

A *service definition* contains all of the information necessary to perform workload management processing. Workload management provides an online panel-driven application for establishing a service definition.

Much of the information you need to establish a service definition is contained in the performance objectives described in “Establishing performance objectives” on page 370. An important piece of information contained in the service definition is the classification rule.

A *classification rule* consists of a work qualifier and a service class. One set of classification rules exists for each service definition. However, you can have multiple service classes for each service definition. Using the workload management-supplied function, you establish the classification rule and define the service class by specifying one or more work qualifiers. The *work qualifier* associates incoming transactions with a particular service class, which represents a group of transactions with similar performance criteria (performance goals and business importance, for example). Work qualifiers can be the subsystem type, the IMS transaction name, the IMS transaction class, the inputting LTERM name, or the user ID. You can use any one of these values, or a combination of them, to assign the service class to a transaction. A *service class* has performance criteria defined for it. After a transaction is assigned a service class, the z/OS Workload Manager processes it according to the performance criteria defined for that service class.

In the following example, the classification rule uses the transaction name as the work qualifier. It assigns all incoming transactions with the name of DEPOSIT to the service class IMSHIGH.

- Work qualifier:
  - Subsystem name: IMS
  - Transaction name: DEPOSIT
- Service class:

- Performance goal: response time of less than a second
- Business importance: 1

The performance goal for the IMSHIGH service class is a response time of less than 1 second. The business importance of 1 gives the DEPOSIT transactions priority, after they are scheduled by IMS, if contention for system resources occurs.

**Recommendation:** Specify multiple WLM service classes, differentiating by business importance and goals. Then classify your IMS transactions in the WLM policy by using the IMS transaction class as the work qualifier for the WLM service that best fits the response time goals of each transaction.

By assigning different WLM service classifications to different IMS transactions, you ensure that the WLM manages the z/OS resources in a manner that gives all IMS transactions the best possible chance of obtaining their response time goals

If you assign a single WLM service class to all of your IMS transactions, the WLM treats all IMS address spaces, including the control region, equally. This means that WLM assigns the same dispatch priority to all of these address spaces. If some of your transactions do not make their response time goals, WLM would not know which address space to give the higher priority to (it should be the IMS Control Region) and would not change the priority for any of them.

## Migrating to workload management

Workload management offers two operating modes: compatibility and goal. Compatibility mode is your present method of performance management, and goal mode is the workload management method.

To facilitate migration, you can establish your service definition while running in compatibility mode. Then, after you are comfortable with your service definition and have completed the steps for migrating to workload management, you can switch to goal mode.

**Related reading:** For more information on setting up service definitions and using workload management, see *z/OS MVS Planning: Workload Management*.

## Interpreting z/OS WLM Change State PB service codes

The WLM Change State Performance Block (PB) service is used to show the current state of the transaction.

The PB service codes are interpreted as follows for IMS:

### State Description

**Active** The transaction is executing an application program.

**Free** This is not reported by WLM.

**Idle** The transaction is waiting for work.

### Waiting - I/O

IMS waiting on I/O (IMS initialed I/O).

### Waiting - Lock

IMS waiting on a lock request.

## Deciding on monitoring activities and techniques

To develop a master plan for monitoring and analyzing performance, you should determine whether dynamic, daily, or detailed monitoring suites your needs best and develop a plan accordingly.

When you develop a master plan for monitoring and analyzing performance, you should establish:

- A master schedule of monitoring activity  
Coordinate monitoring with operations procedures to allow feedback of online events to be incorporated into instructions for daily or detailed data gathering.
- Which tools are to be used for monitoring  
The tools used for data gathering should provide for dynamic monitoring, a daily collection of statistics, and more detailed monitoring.
- The kinds of analysis to be performed  
Document what data is to be extracted from the monitoring output, identifying the source and usage of the data. Although the formatted reports provided by the monitoring tools help organize the volume of data, you should probably design worksheets to assist in data extraction and reduction.
- A list of the personnel that are to be included in any review of the findings  
The results and conclusions from analyzing monitor data should be made known to the user liaison group and to system performance specialists.
- A strategy for implementing changes to the online IMS system design resulting from tuning recommendations  
Coordinate this change implementation strategy with your installation's standards for testing and standards for frequency of production environment changes. Change management is described in Chapter 24, "Modifying the system design," on page 395.

Plan for three broad levels of monitoring activity:

- Dynamic  
Observe the system's operation continuously to discover any serious short-term deviation from performance objectives.  
  
The output from the /DISPLAY or QUERY command is suitable for this level of monitoring, together with end-user feedback. One use of the Resource Measurement Facility™ (RMF™) II is to collect information about processor, channel, and I/O device utilization.  
  
The MTO is an important source of information about the behavior of the online IMS system. An important part of MTO feedback is the set of conditions during an IMS Monitor run. This information can help establish the validity of the monitor data.  
  
With the status information that can be obtained using the /DISPLAY or QUERY command, you can arrange to get a processing status during online execution. The status can include the queue levels, active regions, active terminals, and the number and type of conversational transactions. Such a status can be obtained with the aid of an automated operator program invoked by the MTO. At prearranged milestones in the production cycle—such as before scheduling a message or BMP region, at shutdown of part of the network, or at peak loading—the transaction processing status and measures of system resource levels can be recorded.

To dynamically monitor paging rates, use one of the following methods:



### **Total System Paging Rate**

The RMF Monitor II Paging report gives snapshots of paging activity by sampling interval.

### **Paging Rate by User**

The RMF Monitor II Address Space Resource Data report gives counts of allocated frames and page faults by address space for each RMF measurement interval. The page fault count is cumulative and includes both local address space page faults (not page I/Os) and CSA/LPA.

If the monitoring indicates that IMS is experiencing temporary delays, it might be possible to stop some TSO or batch initiators.

To monitor processor resources, use the following:

### **Total System Utilization**

The RMF CPU Activity report gives WAIT TIME PERCENTAGE for the RMF report interval, which might typically be 10 to 30 minutes.

### **Processor Over commitment**

RMF Monitor II Real Storage/Processor/SRM report gives percentage processor utilization for each RMF sampling interval. It shows a 101% figure if any address space is ready to be dispatched but must wait for processor cycles during the interval.

- **Daily**

Measure and record key system parameters daily. Record both the daily average and the peak period (usually one hour) average. Compare against major performance objectives and look for adverse trends.

This data usually consists of counts of events and gross-level timings. In some cases, the timings are averaged for the entire IMS system, for example, elapsed times for input queuing or program execution.

You can use the IMS system log as input to offline processing to produce statistics on a daily or regular basis. Two utilities, IMS Log Transaction Analysis and IMS Statistical Analysis, are suitable for this level of monitoring, because they impose no additional processing load on the online system.

If the analysis of dynamic RMF Paging Activity reports shows that paging I/O has increased, this must be related to some changes in the workload during the corresponding period. More detailed analysis is usually required to decide on tuning actions.

If analysis of the RMF CPU Activity reports show consistently high CPU utilization during a period of poor response times, examine the elapsed and processor times for transactions by IMS message region. If the elapsed time/processor ratio is significantly higher in the lower-priority regions, raise their dispatching priorities relative to other non-IMS work.

**Related reading:** If all regions are equally affected, see “Minimizing path length” on page 423 for information on reducing path lengths. If appropriate, adjust priorities of competing work to move it below IMS message regions.

To monitor utilization by subsystem for processor over commitment, use the RMF Monitor II Address Space State Data report. This report gives processor units consumed by each address space for each RMF measurement interval. IMS control and message regions can be identified by job name.

- **Detailed**

Periodically collect detailed statistics on system operation for performance analysis against system-oriented objectives and workload profiles.

Data at this level is much more voluminous. It typically contains sequences of events and tabulations. The timings reported are at a detailed level.

At this level of monitoring, special trace tools such as the IMS Monitor and Generalized Trace Facility (GTF) are useful. They collect a detailed sample of the online processing and distinguish between activity in dependent regions, asynchronous processing for terminals and message queues, buffer pool usage, and system data set I/O.

Additional information on using these monitoring tools is included in other topics of this section. The use of monitoring and tools to detect performance problems is explained in “Identifying and correcting performance problems” on page 426.

You can use the following methods for detailed monitoring of paging rates:

#### **Paging Rate by User**

If analysis at the daily level is insufficient, plan to run a GTF trace. If multiple page data sets are being used, private and global paging can be identified. Examine the GTF Detail Trace report to evaluate the impact on IMS of any paging by calculating the elapsed time delays due to page faults, particularly for the control region. The GTF Page Fault Summary report can be used to discover which area of IMS or system code is being affected by page faults. The report also indicates the type of page faults.

Analyze this type of data to help you decide whether to tune real storage usage.

**Related reading:** For other factors to consider when tuning real storage usage, see “Trade-offs between I/O controlled by IMS and paging” on page 420.

#### **Examining NOT-IWAIT Time during Scheduling/Termination**

The IMS Monitor Region Summary report can help you make an initial assessment of dispatching priority problems by examining the Scheduling or DL/I NOT-IWAIT times, that is, the elapsed time not accounted for by IWAIT time. Increases in the NOT-IWAIT times can be caused by:

- Paging delays
- Dispatching for a higher-priority task

If minimal paging is occurring, the portion of elapsed wait time that occurs during the scheduling and termination of a region is fairly consistent from system to system. This elapsed wait time is not accounted for by IWAIT time. This time is recorded in the IMS Monitor Region Summary report and tabulated under the heading NOT-IWAIT TIME. DL/I call NOT-IWAIT times can also include dynamic logging I/O delays. Any delays caused by paging or dispatching for a higher-priority task result in an increase in the NOT-IWAIT times.

If the total mean NOT-IWAIT TIME is excessive, the machine resource is probably inadequate for IMS. If no higher-priority tasks are present, the cause is probably a high paging rate for IMS scheduler code, control blocks, and PSB pool.

The SVC Mapping Summary can assist in determining where an SVC is being issued.

**Related reading:** For more information on reducing path lengths, see “Minimizing path length” on page 423.



---

## Tools for detailed monitoring

Many of the monitoring tools you can use to collect detailed data are also used for general diagnostics. The principal tool provided by IMS for collecting and analyzing data is the IMS Monitor, which allows you to monitor online subsystems.

For a stand-alone IMS DB batch system driven by an SLDS, use the DB Monitor. The DB Monitor can be active during the entire execution of an IMS batch job, or you can stop and restart it from the system console.

You can also use the IMS Performance Analyzer, program isolation and lock traces, and the external trace facility.

### Related concepts:

Chapter 51, “//DFSSTAT reports,” on page 811



IMS Performance Analyzer for z/OS overview

### Related tasks:



Gathering performance-related data (Operations and Automation)

## IMS Monitor

The IMS Monitor collects data while the online IMS subsystem is running. It gathers information for all dispatch events and places it, in the form of IMS Monitor records, in a sequential data set. Use the IMSMON DD statement in the IMS control region JCL to specify the IMS Monitor data set.

IMS adds data to this data set when you activate the IMS Monitor using the /TRACE command. The IMS MTO can start and stop the IMS Monitor to obtain snapshots of the system at any time. However, the IMS Monitor adds to system overhead and generates considerable amounts of data.

### Control of IMS Monitor output

Plan to run the IMS Monitor for short intervals and to control its operation carefully. Shorter intervals also prevent the overall averaging of statistics, so that problems within the system can be more readily identified. The output of the IMS Monitor can be constrained by:

- Type of activity monitored
- Database or partition or area
- Dependent region
- Time interval

### IMS Monitor output data sets

The IMS Monitor output can be either a tape or a DASD data set. Using DASD eliminates the need to have a tape drive allocated to the online system. If you want to use the IMS Monitor frequently, you might find that permanently allocated space for a DASD data set is convenient. One technique is to code DISP=SHR on the IMSMON DD statement so that the reports can be generated as each IMS Monitor run completes.

You must coordinate the report generation with the operator because each activation of the IMS Monitor writes over existing data. Although this overwriting

does not occur for tape data sets, new volumes must be mounted. The volume is rewound, and a mount request is issued each time you start the IMS Monitor.

#### **Recommendations:**

- Do not catalog IMS Monitor data sets. The IMS Monitor can produce multiple output volumes while IMS is running if the data sets are not cataloged.  
If you want IMS to dynamically allocate the IMS Monitor data set, do not include the IMSMON DD statement in the IMS control region JCL.
- Allow IMS to dynamically allocate IMS Monitor tape data sets. A tape drive is not permanently reserved for the control region for dynamically allocated data sets.

### **IMS Monitor traces**

After you establish monitoring requirements, you might be able to restrict the scope of the IMS Monitor activity. Restricting the scope has the advantage of reducing the impact of the IMS Monitor on system throughput. However, do not compromise the collection of useful data.

You can control what specific types of events are traced by using specific keywords on the /TRACE command. For example, you can monitor line activity, scheduling and termination events, activity between application programs and message queues, activity between application programs and databases, or all activity. You can also limit monitoring to:

- Particular databases, partitions, or areas
- Particular dependent regions
- A specified interval of time

### **IMS Monitor reports**

You can obtain reports based on the IMS Monitor output by using the IMS Performance Analyzer for z/OS or the IMS Monitor Report Print utility (DFSUTR20).

The IMS Monitor Report Print utility summarizes and formats the raw data (except Fast Path data) produced by IMS and presents the information in a series of reports. You can suppress the reports pertaining to DL/I calls and tabulated frequency distributions.

The duration of the monitored events is determined by the entries for start and stop of the IMS Monitor. You cannot select a different time period for reporting, because many of the timed events are not captured continuously: only when the IMS Monitor is started and stopped. For this reason, ensure that the IMS Monitor is stopped before taking any action to stop the IMS control region.

#### **Related concepts:**

-  [IMS Monitor \(Database Administration\)](#)
-  [Monitoring the system \(Operations and Automation\)](#)
-  [IMS Performance Analyzer for z/OS overview](#)

#### **Related reference:**

-  [IMS Monitor Report Print utility \(DFSUTR20\) \(System Utilities\)](#)

## z/OS Generalized Trace Facility (GTF)

Use the z/OS generalized trace facility (GTF) to record a wide range of system-level events. The trace activity is controlled from the z/OS system console using the MODIFY command.

Output is spooled to a sequential data set that is used by a generalized formatting utility. You can write exit routines that are called by the formatting utility to edit the trace records and present the data as desired.

The following GTFPARS reports are of special interest for IMS:

- Job Summary report: This report lists the following for each of the IMS message regions and for the control region: SVC counts, loading of modules, and EXCP, SIO, and IO timings.
- Detail Trace report lists the following information:
  - For IMS message regions: Detailed analysis of program management activity. This analysis gives you the ability to investigate event sequences in detail.
  - For the IMS control region: Database I/O that is handled in the control region, and network-related activity.
  - For all IMS regions running in parallel: Contention between regions (for example, contention for access to a program library), and the flow of control from the control region to the MPPs.
- Job Summary report of Master Address Space in z/OS: This report shows system paging activity.
- Detail Trace report of Master Address Space in z/OS: You can use this report for a detailed investigation of delays caused by paging.
- System Summary report: You can use this report for I/O subsystem analysis for all IMS data sets.

### Related concepts:

 z/OS: Generalized trace facility (GTF)

## z/OS Component Trace (CTRACE) Service

IRLM uses the z/OS component trace (CTRACE) service to trace IRLM activity. Because the trace output is in z/OS CTRACE format, you can use IPCS CTRACE format, merge, and locate routines to process the buffer data.

Use the z/OS TRACE CT command to start, stop, or modify an IRLM diagnostic trace. This command can be entered only from the z/OS master console. Entering the commands requires an appropriate level of z/OS authority. IRLM does not support all the options available on the command. You can also start the IRLM tracing by placing TRACE=YES in the IRLM procedure.

### Related tasks:

 Tracing IRLM activity (Operations and Automation)

### Related reference:

 TRACE CT command (Commands)

 z/OS: MVS interactive problem control system (IPCS) CTRACE subcommand

 z/OS: TRACE command

## Obtaining program isolation and lock traces

In an IMS DB/DC or DBCTL environment, you can detect contention for a database segment by examining the output produced by the Program Isolation Trace Report utility (DFSPIRP0).

To get the source data for the utility, issue the /TRACE SET ON PI OPTION ALL command. To stop gathering source data, issue the /TRACE SET OFF PI command. A control statement for the utility can select a start or stop time relative to a specified date.

Tracing the program isolation function can create additional log records. These records contain the enqueue or dequeue requests issued by the program isolation function between sync points as a result of database updates, checkpoints, and message handling events.

The Program Isolation Trace Report utility reports only those events that require wait time. The report identifies the data management block (DMB) name, database control block (DCB) number, the relative byte address (RBA), the program specification block (PSB) name, and the transaction code. The utility sorts all activity by RBA number (shown as ID in the report). The report lists elapsed times for enqueues that required a wait during the trace interval, and totals the number of enqueues for each ID, DCB, and DMB. The requesting PSB or transaction is considered the holding PSB or transaction of the next enqueue waiting for the same segment. A sample report is shown in the following example. In this report, no elapsed wait time is recorded for Fast Path.

```

PROGRAM ISOLATION TRACE REPORT                                PAGE 1
DATE: 08/10/04
TIME: 16:36 TO 16:37
DCB *** REQUESTING *** ELAPSED **** HOLDING *** ID TOTAL DCB TOTAL DMB      TOTAL
DMB NAME NUM  ID  TRAN AND PSB NAMES  TIME    TIME    TRAN AND PSB NAMES ENQ'S ENQ'S ENQ'S
TABLEDBQ 1 0022D020 DE1Q   PROGDE1Q 16:36:54 0:00.061 DE2Q   PROGDE2Q
                                     1
          003BE00C DE2Q   PROGDE2Q 16:36:51 0:00.027 DE1Q   PROGDE1Q
                                     1
          007D901C DE2Q   PROGDE2Q 16:36:34 0:00.036 DE1Q   PROGDE1Q
                                     1
          008EF014 DE2Q   PROGDE2Q 16:36:49 0:00.038 DE1Q   PROGDE1Q
                                     1
          0090401C DE1Q   PROGDE1Q 16:36:50 0:00.072 DE2Q   PROGDE2Q
                                     1
          00A06010 DE2Q   PROGDE2Q 16:36:38 0:00.046 DE1Q   PROGDE1Q
                                     1
          00A1401C DE1Q   PROGDE1Q 16:36:50 0:00.008 DE2Q   PROGDE2Q
                                     1
TABLEDBR 1 002A901C DE2R   PROGDE2R 16:36:40 0:00.034 DE1R   PROGDE1R
                                     1
          0045801C DE2R   PROGDE2R 16:36:41 0:00.043 DE1R   PROGDE1R
                                     1
          0072F024 DE1R   PROGDE1R 16:36:30 0:00.053 DE2R   PROGDE2R
                                     1

```

You can use the File Select and Formatting Print utility to select and print trace table and PI entries in the log records in the following ways:

- Specify an OPTION statement with the PRINT parameter and COND=E and EXITR=DFSERA40 keyword parameters. The output is a report containing the program isolation (PI) trace records formatted in sequential order.
- Select only the log records that contain the trace using the IMS Trace Table Record Format and Print module (DFSERA60), which is an exit routine that receives type X'67FA' log records from the File Select and Formatting Print utility and formats the records on the SYSPRINT data set.

Specify an **OPTION** statement with the **PRINT** parameter and **COND=E** and **EXITR=DFSERA60** keyword parameters. The output is a report containing the PI trace entries, the DL/I trace entries, and the lock trace entries formatted to show these entries in sequential order. For an explanation of the headings, see an assembly listing of the macro **IDLIVSAM TRACENT**.

You can use an output report from the **IMS Trace Table Record Format and Print** module to find out more information about:

- The level of control (**LEV**) column shows read only, share, exclusive control, and single update activity.
- The return code (**RC**) column shows return codes from **DFSFXC10** or the **IRLM**. You can determine whether the caller had to wait for the requested resource, or if the transaction caused a deadlock situation.
- The **PST** post code (**PC**) column shows the cause of the wait. If the entry is **X'60'**, a deadlock occurred.

You can reduce the number of records examined by specifying an additional **OPTION** statement to the **File Select and Formatting Print** utility so that only records confirming deadlock are printed.

IMS automatically resolves deadlock situations by using dynamic backout. But the detection of deadlocks is important so that you can modify your application design to prevent future deadlocks.

The advantage of the **Program Isolation Trace Report** is that it shows where contention for a particular segment or range of segments occurs. The report also shows which transactions are competing within a database. It also shows high wait times that might explain a delay in response time. One way to handle the segment contention might be change the database design to separate some of the fields into an additional segment type.

**Related reference:**

 **Program-Isolation-Trace Report utility (DFSPIRP0)** (Database Utilities)

 **Program Isolation Trace Record Format and Print module (DFSERA40)** (System Utilities)

## IMS Trace facility

You can use the **IMS Trace facility** to write IMS trace tables internally or to an external trace data set.

IMS can write this external trace data set to either **DASD** or **tape**:

- **DASD** data sets can be allocated by **JCL** or can be dynamically allocated.
- **Tape** data sets must be dynamically allocated.

You can also write the trace tables to the **OLDS**, but this could adversely affect **OLDS** performance. The external trace data sets are independent of the **OLDS**, so you can write trace tables to the external trace data sets even if the **OLDS** is unavailable.

To display the status of traces, use the **/DISPLAY TRACE** command. This command can be used to determine the status of the IMS traces in effect and the status of any external trace data sets in use.

**Related concepts:**

➞ Monitoring the system (Operations and Automation)

➞ DL/I trace (Diagnosis)

**Related tasks:**

➞ Gathering performance-related data (Operations and Automation)

**Related reference:**

➞ DFSMDA macro (System Definition)

**Related information:**

➞ DFS2867A (Messages and Codes)

## Analyzing performance in OTMA message routing

You can use the OTMA type-2 command, QUERY OTMATI, to analyze performance in OTMA message routing. You can issue the QUERY OTMATI command in both active and alternate systems regardless of whether you are working in an Extended Recovery Facility (XRF) or remote site recovery (RSR) environment.

If you issue the command QUERY OTMATI without any parameters, IMS displays:

- The workloads on each of the IMS instances in the IMSplex with the target member (tmember) name and the transaction pipe (tpipe) name
- The total number of active messages in the queue as a subset of the total workload
- The length of time that the transaction instance block (TIB) has existed and the correlation ID of the input message

You can use this information to determine if there are problems processing OTMA input messages.

If you include the SHOW parameter with the QUERY OTMATI command, the total number of active messages in the queue is not displayed. Instead, each individual active message is displayed, as filtered by the other specific parameters. For example, to show how many transaction codes have an age of five seconds or more, issue the following command:

```
QUERY OTMATI MSGAG(5) SHOW(TRAN)
```

This command displays the individual TIBs, and shows how many transaction codes have an age of five seconds or more. From this information, you can determine any potential problems in the message queue that can be corrected or avoided, such as storage filling with control blocks, or you can balance the workload of the outbound transactions in the IMSplex without restarting IMS.

---

## Coordinating performance information in an MSC network

The IMS system log for each system participating in Multiple Systems Coupling (MSC) contains only the record of events that take place in that system. The logging of traffic received on links is included.

You can augment the system log documentation that records the checkpoint intervals with the system identifications of all coupled systems. This helps you interpret reports, because you know of transactions that might be present in message queues but are not processed, and you can expect additional transaction

loads from remote sources. In your analysis procedures, include ways of isolating the processing triggered by transactions originating from another system.

To satisfy the need for monitoring with typical activity that includes cross-system processing, coordinate your scheduling of the IMS Monitor and other traces between master terminal operators. The span of the monitoring does not have to be exactly the same, but if it is widely different, the averaging of report summaries can make it harder to interpret the effect of the processing triggered by cross-system messages.

**Related reading:** For more information on interpreting MSC reports, see “IMS Monitor reports for MSC” on page 763.

---

## Monitoring Fast Path systems in DB/DC and DCCTL environments

The major emphasis for monitoring IMS online systems that include message-driven Fast Path application programs is the balance between rapid response and high transaction rates. With Fast Path, performance data is made part of the system log information.

Also, the IMS Monitor can be used to monitor Fast Path activities. You can use the IMS Fast Path Log Analysis utility to generate statistical reports from the system log records. This utility can provide summaries of the Fast Path transaction loads, reports that highlight exceptional response time, and a breakdown of the elapsed time between key events during the time in the system.

The system administration tasks of setting up a monitoring strategy, performance profiles, and analysis procedures should be carried into the Fast Path environment.

**Related reading:** For more information on using either the IMS Monitor or the IMS Fast Path Log Analysis utility, see *IMS Version 11 System Utilities*.

---

## Transaction flow in DB/DC and DCCTL environments

A distinct sequence of events occurs during the processing of a transaction. Message-related processing is asynchronous within IMS, that is, not associated with a dependent region's processing.

Examples of this kind of processing are message traffic, editing, formatting, and recovery-related message enqueueing, any of which can be done concurrently with application program processing for other transactions. Events from application program scheduling to termination are associated with a PST and can be regarded as synchronous.

To get an overall picture of the activity that takes place when an online IMS system is processing a mix of transactions concurrently, see the following figure. The figure shows a sequence of events. Each event is explained in notes that follow the figure.

The unit of work by which most IMS systems are measured is the transaction (or a single conversation iteration, from entering of an input message to receipt of one or more output messages in response). One way of representing the flow of units of work is to compare it to three funnels through which all transactions must pass, as illustrated in the following figure. The events that account for the principal contributions to transaction response time are numbered in the center. The items entered on the left of the diagram are message related, and those on the right are



related to the application program. The arrows trace the flow for an individual transaction. (The diagram does not show the paging element or system checkpoint processing that is distributed through the elapsed times.)




Figure 43. Processing events during transaction flow through IMS

**Notes to the figure:**

1. Wait for Poll.

This is the time between pressing the Enter key and receiving a poll that results in the data being read by the channel program. With several control units on a line, a slight delay is probable, because each control unit is polled in turn until one is ready to send. There can also be delays caused by data being transmitted on the line from (or to) another terminal. Also, IMS communications can be currently processing an input message from a terminal on the line and polling has not recommenced. Some form of hardware monitor is required to measure the time waiting for poll.

2. Data Transfer.



This time includes propagation delay and modem turnarounds for multi-block input messages. You can estimate the data transfer times if the volume of data transmitted is known.

3. Input Message Processing.

The IMS control of the transaction begins when the input message is available in the HIOP. The time the message spends in this pool, in MFS processing, and in being moved to the message queue buffers affects response time. Individual transaction I/O to the Format library affects the message queue. A major factor in determining response time is whether the respective pools are large enough for the current volume of transactions flowing into input queuing. In particular, if the message queue pool is too small, overflow to the message queue data sets occurs.

4. Message Classification.

This is the call to the z/OS WLM to obtain a WLM service classification for the incoming message.

5. Input Queuing.

This is the time spent on the input queue or in the message queue buffers waiting for a message region to become available. In a busy system, this time can become a major portion of the response time. The pattern of programs scheduled into available regions and the region occupancy percentage are important and should be closely monitored.

6. Scheduling.

Because of class scheduling, regions can be idle while transactions are still on the queue. The effects of scheduling parameters can be:

- Termination of scheduling as a result of PSB conflict or message class priorities
- Termination of scheduling as a result of intent conflict
- Extension of scheduling by I/Os to IMS.ACBLIB for intent lists, PSBs, or DMBs
- Pool space failures in either the PSB or DMB pools

7. Init PB Call (Activate Delay Monitoring Environment).

Activate the WLM delay monitoring environment for the message when it is placed into the dependent region. The WLM PB is initialized with the Service Classification and transaction name, message arrival time, program execution start time (current time), userid, and so forth.

8. Program Load.

See event 9

9. Program Initialization.

After scheduling, several kinds of processing events occur before the application program can start:

- Contents supervision for the dependent region
- Location of program libraries and directories to them
- Program fetch from the program library
- Program initialization up to the time of the first DL/I call to the message queue

For monitoring, you can obtain the overall time for the above activities. The number of I/Os should be checked periodically.

10. Message Queue GU.

This is the GU call to the message queue. It is chosen as a measuring point because the event is recorded on the system log and is used as a starting point for iterations of processing when more than one message is serviced at a single scheduling of the program.

11. Program Execution.

The time for program execution, from first message call to the output message insert, is a basic statistic for each individual transaction. It is important to account for that time in terms of the work performed:

- The number of transactions processed per schedule
- The number and type of DL/I calls per transaction
- The number of I/Os per transaction

A useful breakdown of elapsed time into processor time and I/O helps determine which transactions use significant resource.

12. Output Message Insert.

This time begins the asynchronous processing for the output response. The output message requests flow into the funnel to be serviced while the application program is either beginning to process another input message or is performing closeout processing and program termination.

13. Wait for Sync Point.

When an output message is issued by a program, it is enqueued on a temporary destination until the program reaches a synchronization point. For programs specified as MODE=MULT, a long delay in output transmission can occur when the program goes on to process several transactions at one scheduling. None of the previous output messages can be released for transmission until the program ends. If the program fails, all current transactions are backed out. Actually, when the messages are dequeued, LIFO sequence is used, from temporary to permanent destination. With MODE=SNGL, the wait for sync point (at the next GU to the message queue) is normally negligible.

14. Program Termination.

In the case of MODE=MULT, discussed in event 13, the synchronization point occurs at program termination. Any database updates are purged from the database buffer pools, and the waiting output messages are released.

In the MODE=SNGL case, the synchronization point occurs at the previous message queue GU call (usually a GU with a QC status code), and no database commit processing occurs at termination, unless the application program has updated a database after the last message queue GU.

15. WLM Notify Call.

This tells WLM that the application program has ended execution. The PB and current time are passed to WLM.

16. Wait for Selection.

This time is similar to Wait for Poll on input, with one difference—an output message might not have to wait for the completion of a polling cycle if no polling is in progress on the line at the time the output message is enqueued. However, there might be a wait for the duration of data transmission to other terminals on the line. In a busy system, this could account for the majority of time spent on the output queue.

17. Output Message Processing.

This activity is similar to event 3 on page 386.

18. WLM Report Call.

This tells WLM the response is being sent. IMS passes the input message arrival time, the Service Classification, the current time (output message send time).

19. Data Transfer.

This activity is similar to event 2 on page 385.

20. Output Queue Processing.

Output messages that have been sent are dequeued after their receipt has been acknowledged by the terminal. In the case of paged output, the acknowledgment is a consequence of another input or a PA2 entry from the terminal.

**Related concepts:**

“DB/DC environment” on page 6

---

## The IMS Monitor in DB/DC and DCCTL environments

The principal monitoring tool provided by IMS is the IMS Monitor. It is an integral part of the control program in the DB/DC environment. The counterpart of the IMS Monitor in the batch environment is the DB Monitor.

The IMS Monitor collects data while the DB/DC environment is operating. Information in the form of monitor records is gathered for all dispatch events and placed in a sequential data set. The IMS Monitor data set is specified on the IMSMON DD statement in the control region JCL; data is added to the data set when the /TRACE command activates the monitor. The MTO can start and stop the monitor, guided by awareness of the system's status, to obtain several snapshots.

**Related reading:** For more information on interpreting IMS Monitor reports, see Part 5, “Using IMS reports,” on page 707.

---

## Monitoring transaction-level statistics

IMS logs statistical data about transactions to the OLDS in the X'56FA' log record. You can use this log record to collect, analyze, and interpret statistical data for general business purposes.

This statistical information is only a subset of the information that is logged to the X'07' log record. Transactions within a specific scope (such as a unit of recovery) from the X'07' log record are collected. Since the statistical information about the transactions is also logged separately in the X'56FA' log record, you can access this information more easily than extracting it from the X'07' log record.

IMS logs statistical information about the transactions within a unit of recovery (work done by an application program between sync points) in log record X'56FA' for non-message-driven applications, or when either:

- The keyword MODE=SNGL is specified in the TRANSACT macro
- The CMTMODE(SNGL) parameter is specified on a CREATE TRAN or UPDATE TRAN command

IMS logs statistical information about the transactions after each message in log record X'56FA' for message-driven applications when either:

- The keyword MODE=MULT is specified in the TRANSACT macro

- The CMTMODE(MULT) parameter is specified on a CREATE TRAN or UPDATE TRAN command

You can enable or disable this logging on a program basis for non-message driven applications, and on a transaction-by-transaction basis for message-driven applications. Use the following commands to enable or disable this logging on a transaction-by-transaction basis:

- For an existing transaction, issue the UPDATE TRAN or UPDATE TRANDESC command and specify the new TRANSTAT() parameter.
- For a new transaction, issue the CREATE TRAN or CREATE TRANDESC command and specify the new TRANSTAT() parameter.

To enable or disable this logging on a program basis:

- For an existing application program, issue the UPDATE PGM or UPDATE PGMDESC command and specify the new TRANSTAT() parameter.
- For a new application program, issue the CREATE PGM or CREATE PGMDESC command and specify the new TRANSTAT() parameter.

You can enable or disable this logging globally during system definition by specifying a new parameter, TRANSTAT= Y | N, in the Diagnostics Statistics section of the new DFSDFxxx PROCLIB member. This setting applies to any transactions and application programs that are created with the system definition process.

Any transactions or application programs that are created after an IMS cold start using the online change process or the Destination Creation exit routine (DFSINSX0) inherit the TRANSTAT= parameter setting that is specified in the DIAGNOSTICS\_STATISTICS section of the DFSDFxxx PROCLIB member.

To determine whether or not this logging is enabled for a given transaction or application program, issue one of the following type-2 commands:

- QUERY TRAN
- QUERY TRANDESC
- QUERY PGM
- QUERY PGMDESC

---

## Monitoring procedures in a DBCTL environment

This topic explains how to establish monitoring procedures for your DBCTL environment. First, consider that monitoring in a DB/DC environment generally refers to the monitoring of transactions.

The transaction is entered by an end user on a terminal, is processed by the DB/DC environment, and returns a result to the user. Transaction characteristics that are measured include total response time and the number and duration of resource contentions,

A DBCTL environment has no transactions and no terminal end users. It does, however, do work on behalf of CCTL transactions that are entered by CCTL terminal users. DBCTL monitoring provides data about the DBCTL processing that occurs when a CCTL transaction accesses databases in a DBCTL environment. This access is provided by the CCTL making the database resource adapter (DRA) request.

The most typical sequence of DRA requests that represents a CCTL transaction would be:

- A SCHED request to schedule a PSB in the DBCTL environment
- A DL/I request to make database calls
- A sync point request, COMMTERM, to commit the database updates and release the PSB

The DBCTL process that encompasses this request is called a unit of recovery (UOR).

DBCTL provides UOR monitoring data, such as:

- Total time the UOR exists
- Wait time to schedule a PSB
- I/O activity during database calls

This information is similar to, and is often the same as, DB/DC monitoring data. However, in a DBCTL environment, the UOR data represents only a part of the total processing of a CCTL transaction. You must also include CCTL monitoring data to get an overall view of the CCTL transaction performance.

In this topic, the term “transaction” refers to a CCTL transaction. When it applies, UOR is specifically named.

The CCTL administrator must decide what strategy to use to monitor transaction performance. This topic describes some possible strategies and how IMS functions can help.

Several types of monitoring strategies are available. You can:

- Summarize actual workload for the whole of the online execution. This can be continuous or at an agreed-to frequency. Total workload or selected representative transactions can be tracked.
- Take sample snapshots at peak loads and under normal conditions. It is always useful to monitor the peak periods for two reasons:
  - Bottlenecks and response time problems are more pronounced at peak volumes.
  - The current peak load is a good indicator of what the future average will be like.
- Monitor critical transactions or programs that have documented performance criteria.

Plan your monitoring procedures in advance. A procedure should explain the tools to be used, the analysis techniques to be used, the operational extent of those activities, and how often they are to be performed.

You need to:

- Develop performance criteria
- Develop a master plan for monitoring, data gathering, and analysis

## **Establishing performance objectives in a DBCTL environment**

Inherent in the design of your online DBCTL/CCTL system are your performance objectives. Establishing performance objectives is a major task requiring data gathering for the DBCTL/CCTL workload as a whole. After defining the workload

and estimating the resources required, you must reconcile the desired response with the response you consider to be attainable. Monitor the performance of the system to determine if these objectives are being met.

Base performance objectives on:

- Desired, acceptable, and maximum response time
- Average and maximum resource demands (or workload) per transaction
- Predicted and actual transaction volumes

Establishing your performance objectives is an iterative process involving the following activities:

- Defining CCTL user-oriented performance objectives and priorities

These derive from the way a CCTL end user perceives the service provided by the system. These objectives state expectations of response time as seen by the end user at the terminal.

When response time objectives are established, they should not only reflect the time-in-system (the elapsed time from entry of a last input message segment to the first response segment), but also the expected amount of IMS, CCTL, and application program processing. You should consider whether to define your criteria in terms of the average, the 90th percentile, or even worst-case response time. Your choice depends on your installation's audit controls and the nature of the specific transactions.

- Determining how performance against these objectives is measured and reported to users

You can combine UOR data with CCTL transaction data for a complete transaction profile, but keep the following things in mind:

- If a CCTL allows only one UOR per transaction, this one-to-one ratio makes it easy to combine the monitoring data. It is also easier to track, because a CCTL can put a transaction-ID into a SCHED request. DBCTL puts this ID in monitor log records, and it appears in monitor reports.
- If a CCTL allows for multiple UORs within a single transaction, the total UOR data can be obtained by adding the data from each of the UORs.

Besides data from the IMS Monitor, UOR data is also provided to a CCTL upon the completion of a request terminates the UOR. This data can be used in either of the cases above. A CCTL can also perform real-time analysis of UOR statistics related to its transactions.

The same UOR data is also part of the IMS system log, which records the status and end of a UOR.

- Understanding and documenting the current workload

This requires breaking down the total work into categories and, for each category, developing a workload profile (generally estimates) that include:

- Definition of the transaction category (for example, transaction type or group of transactions). Two characteristics of the category are:
  - The transaction profile. Most transactions perform a single function and have an identifiable workload profile. Later in the process, transaction types with common profiles can be amalgamated for convenience.
  - The transaction volume. In situations where the workload to be documented is already operational, a summary of transaction volumes can be obtained. In other cases, volumes are estimated.
- Relative priority of category, including periods during which priority changes
- Resource requirements of the work:



- Physical resources managed by the operating system (real storage, DASD I/O, terminal I/O)
- IMS logical resources managed by the subsystem, such as control blocks, latches, buffers, and number of database resource adapter (DRA) threads
- CCTL resources required

To obtain a base profile of transaction resource demands, you can start DBCTL/CCTL on a dedicated machine and execute a few transactions to accomplish initialization and buffer pool usage. Then start the IMS Monitor and measure a sample of transaction execution.

You can use the base transaction profile to examine the transaction workload to see if there is any way to reduce it. Such design changes result in the greatest impact, occurring before system-wide contention. You can also compare the base profile against the transaction profile in the production environment.

- Translating the resource requirements and volume information obtained into system-oriented objectives for each work category

This includes statements about the transaction rates to be supported (including any peak periods) and the internal response-time profiles to be achieved.

- Confirming that the system-oriented objectives are reasonable

After initializing the system and monitoring its operation, you need to find out if the objectives are reasonable (given the hardware available), based upon the measurements of the workload. If the measurements differ greatly from the estimates used, you must revise the workload documentation and system-oriented objectives accordingly, or tune the system.

## Deciding on monitoring activities and techniques in a DBCTL environment

To develop a master plan for monitoring and analyzing performance, you should determine whether dynamic, daily, or detailed monitoring suites your needs best and develop a plan accordingly.

When you develop a master plan for monitoring and analyzing performance, you should establish:

- A master schedule of monitoring activity  
Coordinate monitoring with operations procedures to allow feedback of online events to be incorporated into instructions for daily or detailed data gathering.
- Which tools are to be used for monitoring  
The tools used for data gathering should provide for dynamic monitoring, a daily collection of statistics, and more detailed monitoring.
- The kinds of analysis to be performed  
Document what data is to be extracted from the monitoring output, identifying the source and usage of the data. Although the formatted reports provided by the monitoring tools help organize the volume of data, you should probably design worksheets to assist in data extraction and reduction.
- A list of the personnel that are to be included in any review of the findings  
The results and conclusions from analyzing monitor data should be made known to the user liaison group and to system performance specialists.
- A strategy for implementing changes to the DBCTL environment design, and to the CCTL system design, resulting from tuning recommendations

Coordinate this change implementation strategy with your installation's standards for testing and standards for frequency of production environment changes. Change management is described in Chapter 24, "Modifying the system design," on page 395.

Plan for three broad levels of monitoring activity:

- Dynamic

Observe the system's operation continuously to discover any serious short-term deviation from performance objectives.

This topic refers to IMS functions and general concepts related to dynamic monitoring. If the CCTL has similar functions (for example DISPLAY commands), they should also be used.

The output from the /DISPLAY or QUERY command is suitable for this level of monitoring, together with end-user feedback. One use of the Resource Measurement Facility (RMF) II is to collect information about processor, channel, and I/O device utilization.

The MTO is an important source of information about the behavior of the online IMS system. An important part of MTO feedback is the set of conditions during an IMS Monitor run. This information can help establish the validity of the monitor data.

With the status information that can be obtained using the /DISPLAY or QUERY command, you can arrange to get a processing status during online execution. The status can include the pool levels and active regions. At prearranged milestones in the production cycle—such as before scheduling a message or BMP region, at shutdown of part of the network, or at peak loading—the transaction processing status and measures of system resource levels can be recorded.

- Daily

Measure and record key system parameters daily. Record both the daily average and the peak period (usually one hour) average. Compare against major performance objectives and look for adverse trends.

The IMS system log can be used as input to offline processing to produce statistics on a daily or regular basis. The 07 and 08 log records contain the UOR data.

- Detailed

Periodically collect detailed statistics on system operation for performance analysis against system-oriented objectives and workload profiles.

Data at this level is much more voluminous. It typically contains sequences of events and tabulations. The timings reported are at a detailed level.

At this level of monitoring, special trace tools such as the IMS Monitor and Generalized Trace Facility (GTF) are useful. They collect a detailed sample of the online processing and distinguish between activity in CCTL threads, buffer pool usage, and system data set I/O.

**Related reading:**

- For information on using the IMS Monitor, see *IMS Version 11 Operations and Automation*.
- For information on how monitoring and tools are used to detect performance problems, see "Identifying and correcting performance problems" on page 426.

System for Generalized Performance Analysis Reporting (GPAR), program number 5798-CPR, is designed as a base for reporting programs (IBM or user-written). It helps summarize sequential activity traces like the IMS system



log, IMS Monitor tapes, and GTF traces. It also contains facilities to print user-tailored graphs from any performance data log or non-VSAM sequential data set. GPAR is a prerequisite for ASAP II, VTAMPARS, and GTFFARS.

---

## Chapter 24. Modifying the system design

This topic summarizes the planning activities needed when changes are necessary in the IMS system design.

The need for changes can result from:

- The integration of new applications
- The staged implementation of an application package
- Modifications to an application design already in production
- Modifications to the network if ETO is not used
- Maintenance of the IMS system at the current level and incorporation of fixes for application program problems
- Tuning of system control parameters
- Redistribution of database data sets
- Storage device changes

All of these changes can affect the IMS system definition and necessitate building a new nucleus and control blocks. The last two changes might only require JCL changes and database reorganization. However, the repercussions can be widespread. With any design change, it is important to control its implementation in order to protect the existing system and the needs of all end users. Implementation of most changes to a production system is best handled by prior testing and a carefully monitored validation period.

The task of administering a change to your IMS online system is likely to involve the following activities:

- Reallocating IMS system data sets and updating execution libraries
- Performing the appropriate system definition tasks
- Reviewing security arrangements, adding changes to RACF user, group, and resource profiles
- Documenting changes in operating procedures
- Documenting changes in recovery procedures
- Being aware of the testing strategy and the timing of cutover into production mode
- Analyzing the performance impact and revising predictions
- Coordinating the use of the IMS Monitor and otherwise monitoring the system before and after implementation
- Collecting comparative performance data

---

### Assessing application changes

Because the design and tuning of an existing production IMS online system involves many people and much machine time, you need to minimize the impact of a change and identify as many potential ramifications as possible.

The following table summarizes how a change in an application program (or your control of the system) can be assessed for impact. It shows the IMS area affected and suggests what other administrative areas of responsibility might be affected.

If you are notified of an application program design change or are tracking a staged implementation, you must extract from the documentation the necessary physical system definition changes. You also need suitable detail to enable you to assess whether your defined environment and operational control are affected.

The first column of the table shows the kind of change and the second column shows the change event you isolate from the documentation. For example, an application system adds a batch message program to generate a report after closedown of a physical part of the network. The following information can be obtained:

**Programs**

Name and PSB, program size, and use of overlay

**Databases**

Already defined, read-only access

**Transactions**

None

**Message formats**

None

**Output**

Name of spooled output procedure

**Terminals**

None

**Network Control**

Whether any change in network availability

**Scheduling**

BMP region: time, frequency, limits

**Exit routines**

None

**Tuning**

N/A

**Security**

No RAS control

Working with this information, you define the APPLCTN macro, perform system definition (CTLBLKS), and tailor the BMP region JCL. Then you add its BMP scheduling instructions to the operations procedure, including a spooled output print step. Alternative strategies can be to perform a MODBLKS system definition and use online change for the cutover to production mode or to use the dynamic resource definition function.

*Table 32. Application change control assessment.*

| Change type | Change event    | Areas affected                                       | Impact and actions                                     |
|-------------|-----------------|------------------------------------------------------|--------------------------------------------------------|
| Programs    | Coding fix      | IMS.PGMLIB                                           | Documentation, log the change                          |
|             | Database access | IMS.PGMLIB IMS.PSBLIB<br>IMS.ACBLIB                  | Monitoring DL/I calls and using online change          |
|             | Addition        | IMS.PGMLIB IMS.PSBLIB<br>IMS.ACBLIB APPLCTN<br>macro | Pools and resource contention, and using online change |

Table 32. Application change control assessment (continued).

| Change type   | Change event     | Areas affected                    | Impact and actions                                                |
|---------------|------------------|-----------------------------------|-------------------------------------------------------------------|
| Databases     | Access           | IMS.ACBLIB                        | Monitoring                                                        |
|               | Structure        | IMS.ACBLIB                        | Database reorganization                                           |
|               | Additions        | JCL, IMS.ACBLIB<br>DATABASE macro | Operations and recovery<br>procedures, and using<br>online change |
| Transactions  | Change content   | TRANSACT macro                    | Recalculate message<br>queues, monitoring                         |
|               | Added            | TRANSACT macro                    | Scheduling algorithm, use<br>of online change                     |
|               | Workload         | Message Queues                    | Recalculate blocking                                              |
|               | Message formats  | IMS.FORMAT                        | Buffers and monitoring,<br>and using online change                |
| Output        | z/OS file        | JCL for devices                   | Operation procedures                                              |
|               | Spool print      | LINEGRP, LINE macros              | Operation procedures                                              |
|               | LTERM            | TERMINAL macro                    | Alternative destinations,<br>operational changes                  |
| Terminals     | VTAM             | System definition macros          | VTAM generation,<br>operating procedures                          |
|               | ETO              | IMS.PROCLIB                       | ETO generation                                                    |
|               | Other            | System definition macros          | JCL for online, operating<br>procedures                           |
|               | Network control  | System definition macros          | VTAM initialization, MTO<br>procedures                            |
|               | Scheduling       | TRANSACT macro                    | Message class algorithms<br>operating procedures                  |
| Exit routines | DL/I             | IMS.SDFSRESL                      | Monitor DL/I calls                                                |
|               | Message edit     | IMS.USERLIB                       | Ensure against abnormal<br>termination                            |
|               | Security         | IMS.USERLIB                       | Ensure security needs are<br>met                                  |
|               | ETO              | IMS.SDFSRESL                      | Monitor ETO                                                       |
|               | User descriptors | IMS.PROCLIB                       |                                                                   |

Table 32. Application change control assessment (continued).

| Change type | Change event           | Areas affected                    | Impact and actions                             |
|-------------|------------------------|-----------------------------------|------------------------------------------------|
| Tuning      | JCL parameters         | IMS.PROCLIB                       | Coordinate operators change                    |
|             | System definition      | IMSCTF macro                      | Monitoring                                     |
|             | Security               | SECURITY macro or EXEC parameters | Coordinate operations                          |
|             | IMS.PROCLIB parameters | DFSDCxxx                          | Coordinate operations                          |
|             | Signon                 | RACF, exit routines               | Validate passwords and coordinate with nucleus |
|             | Terminal               | RACF, exit routines               | Validate passwords and coordinate with nucleus |
|             | Transaction            | RACF, exit routines               | Validate passwords and coordinate with nucleus |
|             | RAS                    | RACF and exit routines            | Validate passwords and coordinate with nucleus |

**Related reading:** For more information about using DRD commands, see “The dynamic resource definition process” in *IMS Version 11 System Definition*.

## Introducing changed applications in an active IMS system

When introducing a changed application program into an active IMS system, the effects of quick reschedule might leave an older, unchanged copy of the application program loaded in an MPP, IFP, or JMP dependent region.

To ensure that a refreshed copy of a changed application program is loaded into an MPP, IFP, or JMP dependent region, follow these steps:

1. Issue the /DISPLAY PGM command to identify associated transactions that are using the application program that you are changing.
2. Issue the /DISPLAY TRAN command to identify the scheduling class for each transaction identified in step 1.
3. Issue the /DISPLAY ACT command to identify the dependent regions that are capable of processing the scheduling classes identified in step 2.
4. Stop the application program that you are changing by issuing the /STOP PGM command.
5. Refresh the program.
6. Stop and restart all dependent regions identified in step 3. Stopping and restarting these regions one at a time allows the remaining active regions to process transactions that are not associated with the stopped program during this process.
7. Start the changed application program.

## Resizing the inactive ACB library online

Occasionally, the size of the ACB library needs to be adjusted. The methods of bringing the change online differ depending on how the data set is allocated; with JCL or dynamically using the DFSMDA macro.

When JCL is used for allocation, IMS must be brought down and restarted to recognize the changes. To avoid having to take IMS down to perform this task, use

dynamic allocation for both the active and inactive ACBLIB data sets and then perform an online change process to bring the resized ACB libraries online.

### Example

Due to rapid application development cycle, an ACB library might need to be enlarged frequently. If both the active and inactive ACB libraries are dynamically allocated using the DFSMDA macro, resizing of the ACBLIBs will not require an IMS outage.

## Migrating static allocation of the ACB library to dynamic allocation

Create DFSMDA members to dynamically allocate the inactive ACB library and then restart IMS to use the new DFSMDA members.

To migrate to using the DFSMDA macro for dynamic allocation of ACBLIBs, complete the following steps:

1. Create DFSMDA members for the ACBLIBA and ACBLIBB data sets.
2. Remove the IMSACBA and IMSACBB DD statements from the IMS and DL/I JCL procedures.
3. Add a DD statement that points to the data set where the new DFSMDA members reside. The DD statement can be in either the IMS STEPLIB concatenation or IMSDALIB concatenation.
4. Stop IMS and then restart it with the DFSMDA members.

Now, you can perform the following tasks without stopping and restarting IMS:

- Increase the size of the inactive ACB library data sets.
- Correct errors in the inactive ACB library.
- Add additional data sets to the inactive ACBLIB concatenation.

---

## Planning for system definition changes

The IMS online system design reflected in the macro specifications is subject to change. This design is rarely a one-time execution that occurs at initial installation.

The allocation of real resources and terminal connections change, if not because of hardware changes, then because of application workload and design changes. The ongoing task of monitoring for performance is likely to modify elements of the online system design.

## Controlling system definition processing

When planning a change to system definition, you usually have to coordinate a series of macro changes. The IMS system definition stage 1 input is suggested as a point of control.

Review changes to the macros for the accuracy of the parameter values and add adequate explanation of how the value was derived. You should have a control document that records the nature of the proposed change, the reason for it, and suitable authorization. You can then complete the record by entering the action taken, or planned, for each item.

You must also coordinate the timing of the system definition stage 2 processing. The new version of the control program, with any required JCL changes, must be tied to an operational cutover.

Re-specifying a parameter on one or more macro statements does not require a full system definition or the processing of a complete stage 2 jobstream. You can make many tuning changes by overriding parameter values with others specified in the JCL statements. You can delay making the new values a permanent part of the control program until a more extensive change occurs in the macro content. Consider that all such overrides need to be coordinated with PROCLIB contents and operator instructions.

## Determining the type of system definition required

Your requirement to perform a full system definition or a partial one depends on the macro statements altered. However, the modification of a particular keyword might be the only cause for full generation.

Because the amount of processing performed during a stage 2 execution is significant, you should plan to take advantage, wherever possible, of economics in the type of generation selected.

**Related reading:** For a summary of types of system definition to choose, see *IMS Version 11 System Definition*.

## Changing the network definition in DB/DC and DCCTL environments

You can frequently change terminals that are connected to an IMS online system that have attendant buffer storage tuning. If you are introducing static VTAM terminals into the system, you must regenerate the nucleus. (An online system definition is necessary if no VTAM support was previously defined.) At the same time, you can delete support for specific terminal types.

If you are introducing VTAM terminals and using ETO, it is not necessary to regenerate the nucleus. When you add an ETO terminal supported by the existing descriptors, IMS uses the appropriate ETO logon descriptor to define the terminal.

When considering deleting terminals, leaving the definitions in place might be more efficient if the terminal is physically removed and messages cannot be routed to that LTERM. You can make the deletions in a later definition change.

If you are adding MSC for the first time or adding physical links, you need an online system definition. If, because of a change in the number of IMS systems, you are deleting a requirement for main storage-to-main storage or MSC VTAM connections, you only need to regenerate the nucleus. If you are changing link or system identification names, bind the regenerated control blocks.

### Related tasks:



Creating logon descriptors (Communications and Connections)



Extended Terminal Option (ETO) (Communications and Connections)

## Coordinating system definition and current security definitions

When you are altering the specifications that control security options, you must coordinate the execution of system definition and changes to operating procedures. If you have used IMSGEN and COMM macros to specify security options, you can use the SECURITY macro or initialization EXEC parameters to override any of the options.

The keyword SECLVL on the SECURITY macro determines what overrides are available to the MTO during restart. If you change any of the values for this keyword but not those for TYPE, you only need to perform a CTLBLKS generation and bind. The operational impact is described in detail in “Security considerations for the master terminal” on page 315.

Remember that any change that requires a new or re-bound nucleus must be built into the control program and the current security options.

---

## Making system tuning changes

Through ongoing monitoring and awareness of end-user feedback, you should be aware of potential tuning activities.

The tools and strategies for performance-related tuning activities are described in Chapter 23, “Collecting and interpreting IMS monitoring data,” on page 369, and Chapter 25, “Tuning the system,” on page 407. These sections assume that performance administration is an iterative task made up of:

### Data collection

Using monitors and traces

### Data reduction

Using report processors and manual procedures

### Data analysis

Finding problem indicators

### Change implementation

Proposing solutions and modifying the system

This plan should include a tuning and maintenance section. You can combine items from this section with other required changes to save additional system definition processing and reduce the frequency of changes to the operating procedures. An overall design change log enables you to keep track of the proposals that are analyzed and approved, as well as those under consideration for future implementation. A control document like this also assists in detecting conflicting changes.

The documentation for a change that affects the IMS online system design should include the reasons for the change and a checklist of what parameters or components in the system require modification. Ideally, the documentation should have sections on operations, database maintenance, recovery, security, and monitoring arrangements.

## Resource utilization changes

Certain changes can promote a more efficient use of virtual storage and minimize I/O for an existing system.



These changes include:

- Re-specifying EXEC parameters for the control region
- Allocating DASD storage for message queues
- Database reorganization and data set distribution
- Putting ACBs into 64-bit storage

## Application and database design changes

As a result of a performance study for an application package, you might want to make changes to the online system design.

These changes can range from a simple update of application programs in IMS.PGMLIB, to redefining DBDs and ACBs in IMS.DBDLIB and IMS.ACBLIB. If application programs are changed in IMS.PGMLIB, the message processing regions must be stopped and restarted to store the correct directory entry in the BLDL list.

## Communication design changes in DB/DC and DCCTL environments

If you detect a bottleneck whose symptom is the transmission rate between a communication device and IMS, the response falls into the same planning as for network definition changes.

---

## Managing online system definition changes

A request to add an application or to modify the current set of programs, transactions, and database usage need not force a complete system definition and an IMS restart, with possible interruption for the online users. You can examine the request to see if an online change or dynamic resource definition (DRD) process can satisfy the request.

If the request does not involve changes to the IMS network or the use of static terminals as defined in the current IMS system definition, you can use the online change process to make the changes while the IMS system is executing online. Manage the preparation of the IMS system data sets as described under “IMS System Data Sets for Online Change” in *IMS Version 11 System Definition*. It is also important to assess the impact to the terminal users and the online system operation.

### Related reading:

- For a detailed description of using DRD to make system definition changes, see *IMS Version 11 System Definition*.
- For an overview of the online change function, see “Overview of the local online change function” on page 438.

## Deciding if system modifications can use online change

A list of the changes that can be serviced by online change is given in the following table.

The list is based on the effect on the stage 1 system definition macros, because the stage 1 input is regarded as the control point.

Table 33. System definition resource modifications allowed for online change

| System definition macro | Permissible online changes                                                        |
|-------------------------|-----------------------------------------------------------------------------------|
| APPLCTN                 | Add a PSB (application) and its attributes<br>Change attributes<br>Delete a PSB   |
| DATABASE                | Add a database and its attributes<br>Change attributes<br>Delete a database       |
| RTCODE                  | Add a routing code and inquiry attributes<br>Delete a routing code                |
| TRANSACT                | Add a transaction and its attributes<br>Change attributes<br>Delete a transaction |

The changes listed in the table require a MODBLKS generation. The corresponding security changes require that you also update the RACF security profiles to reflect any new IMS resources needing security.

Also, the system definition preprocessor can be a useful part of your preparation. When adding or changing resource names, the preprocessor can detect any invalid names or duplicate names and help ensure a successful system definition run. If the LGEN subparameter is specified in the SYSTEM= keyword on the IMSCTRL source statement, stage 1 processing occurs during execution of the preprocessor. Some limitations to the modifications are explained in “Planning considerations for online change.” You cannot make modifications that affect the terminal network.

## Planning considerations for online change

When assessing whether a set of changes can be handled with an online change, you must take into account several limitations. In general, the checking performed by the stage 1 processing does not tell you if you have made a change that cannot be implemented online.

You need to consider effects of the following:

- APPLCTN macro

If the message class is assigned as part of the PGMTYPE parameter, that class cannot exceed the maximum number of message classes currently defined for the system.

If the transaction is designated for Fast Path, Fast Path must be active in the system.

Routing a transaction to another IMS system requires that the system name (SYSID parameter) and the use of MSC are not newly defined for the MODBLKS generation.

Although you can make changes to the RESIDENT and DOPT characteristics, PSBs defined as RESIDENT operate as nonresident until after the next restart, because the action of making PSBs resident takes place at IMS system initialization time.

Changing the scheduling attribute to a resident PSB causes that PSB to become nonresident until the next IMS restart.

If a BMP program becomes a message processing program, the transaction characteristics defined in the TRANSACT macro that control message scheduling do not take effect until after the next restart. However, the MTO can use the

/ASSIGN command to specify appropriate message class and processing priorities for the particular transaction. The transaction then becomes eligible for normal message scheduling.

You can use the UPDATE PGM command to update both status and attributes. Some changes you make to the APPLCTN macro characteristics are not implemented as part of the online change processing and only become effective at the next cold start of the IMS online system. They are as follows:

#### TRANSTAT

- DATABASE macro in DB/DC and DBCTL environments

Although the RESIDENT characteristic can be added, the process of making DMBs associated with the database resident does not take effect until after the next restart of the IMS online system.

Changes to the ACCESS parameter are not part of online change; this change can be handled with the /START DATABASE command. Alternately, you can use the UPDATE DB START(ACCESS) command.

You cannot include any kind of change to MSDBs.

- RTCODE macro

The addition of this macro statement, or changes to its specification, is only allowed if Fast Path is active in the system. Make sure that the existing Fast Path User Input Edit routine is able to handle any added routing codes.

- TRANSACT macro

You can control several of the characteristics specified by this macro using such commands as /ASSIGN, /MSASSIGN, /START, and /STOP. You can also use UPDATE TRAN to update both status and attributes. Some changes you make to the TRANSACT macro characteristics are not implemented as part of the online change processing and only become effective at the next cold start of the IMS online system. They are as follows:

PRTY  
PROCLIM  
PARLIM  
SEGNO  
SEGSIZE  
SYSID  
TRANSTAT

Transactions designated as Fast Path potential need Fast Path to be active in the current system.

Routing a transaction to another IMS system requires that MSC facilities be active in the current system. You cannot introduce a system name (SYSID parameter) that was not previously defined in the current system.

Edit exit routines specified for the transaction must already be part of the current IMS online system.

- Page fixing

No additional page fixing is done for added control blocks until the next restart of IMS.

- EMHB size

If you use online change to add or change a transaction-specific EMHB size, ensure that the new EMHB size is not larger than the EPSESRT size. The EPSESRT size is determined only during initialization.

During normal transaction processing, IMS checks the size of the input message against the EMHB length and the EPSESRT length. If the input message exceeds either the EMHB length or the EPSESRT length, it is rejected with message DFS0444.

---

## Performing capacity planning

Other topics discuss short- or medium-range planning for changes in the processing requirements of your IMS online design. However, you might want to examine the long-term adequacy of your design.

For example, if you extrapolate trends in workload, will enough computing power be available?

You can use:

- Individual transaction profiles
- An overall processing profile
- An estimate of percentage processor utilization

As a result of your performance monitoring, you can establish transaction profiles. One part of these profiles can be a trend analysis. You must determine mean and peak arrival rates: currently predicted, sample actuals, and future expectations. The overall processing summary can be taken from the IMS Monitor Run Profile report.

One way to investigate limiting capacity is to extrapolate the volume of transactions. You can establish a trend from percentage processor usage for several monitor points. A common observation from performance measurements indicates that the increase in machine cycles is nearly linear with the increase in transactions per second on systems dedicated to IMS. If you plot several points, showing increasing transaction loads against percentage utilization, you can extrapolate to find your limiting number of transactions as the percentage utilization approaches 100%.

You can also see if a predicted maximum workload corresponds to a utilization that is below a practical upper limit. The workload is estimated by weighing the current overall profile with additional transactions.

In making such predictions, you should have adequate virtual storage and an I/O configuration sufficient to handle the increased transaction load. A careful consideration of remaining available I/O and virtual storage resources is necessary, because contention for these elements has a more severe impact on performance than the level of processor utilization.

If you are managing a large system running under z/OS, you might want to reduce your requirements for CSA. You can plan to use the DL/I separate address space and, if necessary, local storage for IRLM for lock management. Using your projected workload, you might still want to estimate virtual storage with various configurations.



---

## Chapter 25. Tuning the system

You can make changes to parameters that allow the IMS online system to execute more efficiently.

Specifically, these changes to parameters allow IMS to:

- Attain performance criteria
- Efficiently use resources

Chapter 23, “Collecting and interpreting IMS monitoring data,” on page 369 explains setting up performance objectives and monitoring, which are prerequisites to tuning activities for the IMS online system. Terms defined in that section are not generally explained again. Your data for analysis probably depends on the reporting capabilities of one or more of the following tools:

- The Log Transaction Analysis utility (DFSILTA0)
- The Statistical Analysis utility (DFSISTS0)
- The IMS Monitor
- Report programs executing under GPAR: IMSASAP II
- Resource Monitoring Facility (RMF II)

This section presents a general strategy for tuning the IMS online system. After some explanation of performance management, the topics in this section approach tuning in two stages:

- Initializing z/OS and IMS parameters
- Identifying and correcting performance problems

For each activity, several major areas of concern are identified. You should examine the potential tuning activities in each area. This is an outline strategy and not a substitute for more formal performance studies. However, this information should prepare you to address most major problems.

Tuning the system involves initializing z/OS and subsystem parameters and identifying and correcting problems as they relate to tuning. These two topics can be characterized as guidance to reduce the chance of initial problems and a methodology for ongoing problem resolution. However, you should begin by understanding how they relate to the other aspects of the performance management process.

Tuning actions and system redefinition are two aspects of IMS performance management that are part of a continuous cycle of activities, as illustrated in the following figure. Interpreting system performance includes both design decisions and prediction.




Figure 44. Activity cycle for IMS performance management

**Related reading:**

- For information about monitoring and tuning MSC links, see *IMS Version 11 Communications and Connections*.
- For information about tuning Fast Path, see *IMS Version 11 Database Administration*.
- For information about tuning a shared-queues environment, see “Tuning for performance in a shared-queues environment” on page 232.

## Managing performance

IMS performance management is made up of a set of activities, designed to achieve a set of goals.

The goals consist of:

1. Establish performance objectives for each subsystem (IMS, TSO, Batch) and their relative priorities.
2. Initialize MVS and subsystem parameters with a view to meeting the defined performance objectives.
3. Monitor the operation of the total system to confirm that performance objectives are being met.

The three broad levels of monitoring are:

- Dynamic monitoring to discover if performance deviates from established objectives
- Daily monitoring of key performance objectives in order to observe trends and early indicators of possible performance problems
- Detailed monitoring performed periodically to review performance and decide on tuning requirements

For more information about these levels of monitoring, see Chapter 23, “Collecting and interpreting IMS monitoring data,” on page 369.

4. Identify a problem that has been monitored and correct it with minimum disruption of normal service to users.
5. Perform capacity planning or trend analysis to ensure that projected loads will not cause performance problems in the foreseeable future.

Establishing performance objectives and monitoring operations are addressed in “Monitoring the system” on page 349. Performing capacity planning or a trend analysis is described in “Performing capacity planning” on page 405 and in “Using the IMS Performance Analyzer for z/OS” in *IMS Version 11 Operations and Automation*.

## Change management

Performance management is really the management of change. Performance is stable if few changes occur in an IMS system.

More common is a continually changing environment where the IMS system must be adjusted to the addition of new application programs and respond to changes in:

- Application design
- Transaction volumes
- Database volumes
- Workload mix
- User population
- User priorities
- Volumes of non-IMS work
- Hardware configuration

IMS tuning attempts to allocate resources for efficient transaction processing.

You can use various techniques to tune the system during operation. For example, you can:

- Start new message regions
- Change transaction PARLIM, PROCLIM, or class assignments
- Alter message region class assignments

These activities involve ensuring that sufficient IMS resources, such as regions, pool space, or control block areas, are provided. The IMS configuration needs to be matched to the installation's physical resource constraints. Typically, the use of virtual storage is exchanged for some I/O activity controlled by IMS, and contention is reduced for data stored on DASD devices. Tuning should not be crisis management but a continual compensation for changes or workload trends.

A systematic approach to tuning is included in “Identifying and correcting performance problems” on page 426.

## Design variables

The performance of an IMS system depends on the interaction of the workload with many design variables.

Performance is affected by:

- IMS system definition



- IMS execution-time options and parameters
- Application design
- Database design
- Operating system configuration
- Hardware configuration

These areas are addressed in “Initializing z/OS and IMS parameters for tuning.” Design aspects also include allowing for design and tuning changes decided upon as a result of monitoring the system in operation.

## Types of prediction

Two aspects of prediction exist: implicit and explicit prediction.

### Implicit

When an alteration, classified as a tuning change, is made to the online IMS design, a performance improvement is implied.

### Explicit

When actions are taken to make predictions, analytical or simulation models are used, and capacity planning is performed.

---

## Initializing z/OS and IMS parameters for tuning

The IMS and z/OS parameter initialization steps taken reflect the actions that can have a performance impact.

These parameters must be set while considering other aspects. For example, if design or function requirements that contradict or override these recommendations exist, you must evaluate the alternatives.

Initializing z/OS and IMS parameters involves the activities explained in the following topics.

## Assigning z/OS dispatching priorities

Choose a coordinated dispatching priority scheme for the total system. The scheme must reflect the relative worth of any subsystems that are to execute concurrently.

The IMS control region operates in protect key 7 as a system task. This reduces the number of executable instructions for some supervisory functions.

### Job dispatching priority

The job dispatching priority dictates the sequence jobs are given to available machine cycles. Any job with a dispatching priority higher than the control region causes interference. Although VTAM or TCAM normally runs at high priority (as do JES and the z/OS Master Scheduler), you should review the position of other work (principally TSO and batch) in relation to the IMS control region and the dependent regions.

### Dependent region dispatching priority

It is important that both the control region and the dependent regions obtain priority processing. The parallel DL/I function in IMS enables most of the DL/I call processing to be performed in the dependent regions. If the assigned dispatching priority is too low, the promptness of service to application programs

declines. Usually, you set the BMP region priority lower than MPP priorities. The BMP priorities should, however, immediately follow the MPPs, because they contend for the same control region resources, such as program isolation enqueues. Normally, TSO “first period” work is placed below the IMS control region; it might be above the dependent regions. TSO “second period” work and batch are normally placed below the IMS dependent regions.

## Setting a performance group number under z/OS

You can assign a unique performance group number to each IMS region. In this way, you can obtain Resource Measurement Facility (RMF) report statistics separately for each region.

If the DL/I separate address space option is used, you might decide to include both the control and DL/I separate address spaces in the same RMF performance group.

## Setting priorities in a DBCTL environment

For a DBCTL environment, the CCTL region should have a higher priority than the DBCTL region, which should have a higher priority than any batch processing regions.

CICS might run at a higher priority than the DBCTL region. Adjust the priorities of the DBCTL and CICS regions depending on the workload and resources available to your system.

## Dispatching priorities

Dispatching priorities are set by the z/OS Workload Manager (WLM) and are based in part on the relative importance you assign in the WLM ISPF panels to each IMS address space.

**Related reading:** For more information on assigning WLM importance, see “Workload Manager and IMS” on page 372.

The following figure shows an example WLM ISPF display of the relative dispatching priorities assigned to all address spaces known to WLM. The dispatching priority of each address space appears as a hexadecimal value in column DP, with X'FF' being the highest priority. Address spaces with a lower value have a lower priority.

| NP | JOBNAME  | StepName | ProcStep | JobID    | Owner   | C  | Pos | DP   | Real | Paging   | 00000100 |
|----|----------|----------|----------|----------|---------|----|-----|------|------|----------|----------|
|    | ALLOCAS  | ALLOCAS  |          |          |         | NS | FF  | 806  | 0.00 | 00000110 |          |
|    | IOSAS    | IOSAS    | IEFPROC  |          |         | NS | FF  | 256  | 0.00 | 00000120 |          |
|    | IXGLOGR  | IXGLOGR  | IEFPROC  |          |         | NS | FD  | 803  | 0.00 | 00000130 |          |
|    | SMS      | SMS      | IEFPROC  |          |         | NS | FE  | 329  | 0.00 | 00000140 |          |
|    | SMF      | SMF      | IEFPROC  |          |         | NS | FF  | 363  | 0.00 | 00000150 |          |
|    | JES2AUX  | JES2AUX  |          |          |         | NS | F9  | 120  | 0.00 | 00000160 |          |
|    | IRL2194  | IRL2194  | IRLM     | STC00312 | ++++++  | NS | FD  | 283  | 0.00 | 00000170 |          |
|    | NETMONC4 | NETMONC4 | SYSTEM   | STC00005 | ++++++  | NS | FB  | 240  | 0.00 | 00000180 |          |
|    | JES2     | JES2     | IEFPROC  |          |         | NS | F9  | 3095 | 0.00 | 00000190 |          |
|    | VTAM94   | VTAM94   | IEFPROC  | STC00002 | ++++++  | NS | FD  | 2492 | 0.00 | 00000191 |          |
|    | SOFIMSP  | SOFIMSP  |          | STC00006 | ++++++  | LO | FF  | 226  | 0.00 | 00000192 |          |
|    | MPC4A04  | REGION   |          | JOB00323 | USRT008 | A  | NS  | F3   | 385  | 0.00     | 00000200 |
|    | MPC4A05  | REGION   |          | JOB00324 | USRT008 | A  | NS  | F3   | 385  | 0.00     | 00000300 |
|    | MPC4A06  | REGION   |          | JOB00325 | USRT008 | A  | NS  | F3   | 385  | 0.00     | 00000400 |
|    | MPC4A07  | REGION   |          | JOB00326 | USRT008 | A  | NS  | F3   | 385  | 0.00     | 00000500 |
|    | MPC4A08  | REGION   |          | JOB00327 | USRT008 | A  | NS  | F3   | 385  | 0.00     | 00000600 |
|    | SCIC94   | SCIC94   | IEFPROC  | STC00313 | ++++++  | NS | FD  | 460  | 0.00 | 00000700 |          |
|    | OMC94    | OMC94    | IEFPROC  | STC00314 | ++++++  | NS | FD  | 442  | 0.00 | 00000800 |          |
|    | CQSC94   | CQSC94   | IEFPROC  | STC00315 | ++++++  | NS | F7  | 1052 | 0.00 | 00000900 |          |
|    | RMC94    | RMC94    | IEFPROC  | STC00316 | ++++++  | NS | FD  | 507  | 0.00 | 00001000 |          |
|    | I91CTL4  | I91CTL4  | IEFPROC  | STC00317 | ++++++  | NS | F7  | 12T  | 0.00 | 00001100 |          |
|    | I91DLSC4 | I91DLSC4 | IEFPROC  | STC00318 | ++++++  | NS | F5  | 27T  | 0.00 | 00001200 |          |
|    | I91DBRC4 | I91DBRC4 | IEFPROC  | STC00319 | ++++++  | NS | FD  | 617  | 0.00 | 00001300 |          |
|    | BPXOINIT | BPXOINIT | BPXOINIT |          |         | IN | FF  | 162  | 0.00 | 00001400 |          |
|    |          |          |          |          |         |    |     |      |      | 00001500 |          |

Figure 45. Relative dispatching priority (column DP) displayed in the WLM ISPF panel

## Choosing IMS options for performance

Performance implications are associated with many parameters. Some parameters are chosen for system definition, some for execution-time options, and some are an integral part of design decisions.

The following list discuss many of the options.

- Optimizing IMS and application module loading

Where possible, unless availability of virtual storage is a constraint, use the pageable link pack area (PLPA) for IMS reentrant modules, region and program controllers, frequently used high-level language modules, any reentrant application code, and the overlay supervisor (if used).

Use the execution-time parameter SRCH=1 in the IMS procedure to cause the job pack area and link pack area to be searched before STEPLIB or JOBLIB.

- Program library considerations

Place a program library containing the production programs first in the STEPLIB concatenation for the message regions. The sequence of programs in IMS.PGMLIB should be in descending order based on frequency of use.

The order of the system search for program libraries can account for some inefficiency in program load. STEPLIB or JOBLIB is the first program library searched. If one of these is not used, IMS.PGMLIB should be first in the LNKSTnn member of SYS1.PARMLIB. The placement of IMS.PGMLIB is far more important than IMS.SDFSRESL, because its contents are in continual demand for application program scheduling rather than being in demand only at system initialization. The search and load times are significant because they become a performance penalty each time the program is scheduled. You should block programs to full track size in the library. The DC option of the linkage editor, used for downward compatibility, causes program block sizes to be 1024 bytes. You should bind again without this option.

- Dependent region BLDL list

A list of entries is maintained in the dependent region containing the directory index for programs. It is maintained on the most active, most recently used basis, and reduces the I/O to the program directory. Programs with entries in this list have a lower schedule-to-first DL/I call elapsed time than infrequently used programs. You can override the default of 20 entries for the message processing region by using the DBLDL parameter in the EXEC statement of the DFSMPR procedure. The maximum number of entries for the message processing region is 9999.

For a region that is used to test new or changed programs, set the DBLDL parameter to 0, ensuring that the most current version of the program is loaded for each execution. Specifying the DOPT parameter disables quick reschedule.

- Application control block placement

To reduce the amount of read I/O to the ACBLIB data set, you can load the ACB members into 64-bit storage by specifying ACBIN64=ggg in the DATABASE section of the DFSDFxxx PROCLIB member. The value specified (ggg) is the amount of 64-bit storage (in gigabytes) to be allocated for PSB and DMB ACB members.

When the ACB members in 64-bit storage function is enabled, the ACB members are retrieved from 64-bit storage at application scheduling time instead of from the ACBLIB data set.

- Application program control

Specify MODE=SNGL on the TRANSACT macro statement to reduce message queue I/O activity, and response time.

For the batch message program execution controlled by a time limit, select the IMSBATCH procedure parameter STIMER=2. This causes the STIMER/TTIMER macro to be issued only once per schedule, rather than once per DL/I call with STIMER=1. Similarly, use the value 2 for the positional parameter STIMER for the DFSMPR procedure when controlling message processing regions.

For the IMSBATCH procedure, use parameter values of SPIE=0 and TEST=0 for production programs to eliminate unwanted SVCs. The equivalent positional parameters for the DFSMPR procedure are SPIE and VALCK.

With PL/I, use the latest release of the optimizing compiler to reduce initialization and termination overhead.

- Setting checkpoint frequency

Adjust the checkpoint frequency so that checkpoints are not taken more frequently than once every 10 to 20 minutes. You control this frequency with the CPLOG keyword on the IMSCTF system definition macro.

- Message format options

Specify FILL=NULL or FILL=PT in DOFs that have many DFLDs to minimize the transmission of blank characters.

- Setting queuing options

If you do not want priority processing, you can use the IOS queuing option to set priority processing off. The IOS option causes IMS to process all jobs in the queue on a “first-in, first-out” basis.

- Page fixing IMS resources

If you want to ensure that the IMS virtual storage is always backed by real storage, you can specify a list of page-fix requests in the member DFSFIXxx in IMS.PROCLIB. The page fix is done during the IMS control region initialization. Long-term page fixing for the queue manager buffers can be requested using the EXVR parameter in the EXEC statement of the control region JCL.

You tune buffer pools to make them large enough to reduce or minimize I/O activity without driving up the system paging rate or constraining virtual storage. Considerations for page fixing buffer pools are given in “Page fixing the IMS buffer pools” on page 420.

- Defining IMS resources for DREF storage in DB/DC and DCCTL environments  
You can request that the IMS virtual storage never migrate to auxiliary storage. By specifying DREF requests in the DFSDRFxx member in the IMS PROCLIB member, you ensure that virtual storage does not move beyond expanded storage. If expanded storage is not available, the storage is page fixed.

**Related reading:** For more information about DREF requests, see *IMS Version 11 System Definition*.

- MVS page fix considerations  
MVS modules that are frequently referred to by the IMS system should be page fixed at initial program load.  
Page fixing certain modules also saves a “Page fix and Wait” SVC for each occurrence of timer facilities. Place these modules close together to minimize the number of pages in the pageable link pack area (PLPA).

**Related reading:** For more information on MVS modules, see *z/OS MVS Initialization and Tuning Guide*.

## Avoiding contention for IMS resources (excluding buffer pools)

The interleaved processing of concurrent regions is a major area of contention for IMS resources. A set of regions processing a mix of transactions can overlap processing and use the advantage of multitasking over serially ordered execution.

The following list identifies several areas concerned with the pattern of scheduling and region efficiency:

- Using class scheduling in DB/DC and DCCTL environments  
The following points affect your decisions about using class scheduling:
  - Use classes to separate low-priority and long-running transactions from high-priority work.
  - Arrange your class structure so that any high-priority class work has the opportunity of being scheduled in more than one region, unless doing so is undesirable for some special reasons, such as program isolation conflicts. Multi-server systems generally produce lower waiting times than single-server systems.
  - To prevent message regions from being idle while transactions are waiting on the input queue, assign additional (lower-priority) classes to a message region below the primary class or classes assigned to the region. Avoid regions that only service one class, unless you are sure that work of that class will always be waiting to be processed.
  - For very high-volume transactions, dedicate a message region and classify the program as a wait-for-input program to reduce input queuing time and to eliminate scheduling and program load.
- Processing limit counts in DB/DC and DCCTL environments  
Review your transaction class and priority scheme. Set the PROCLIM parameter on the TRANSACT macro to ensure that no transaction type can monopolize a message region if other transactions are waiting and are eligible for processing in that region. Only use a PROCLIM value higher than 5 if the response times of contending transactions are unimportant. In that case, set PROCLIM=1 for the

lower-priority transactions. With wait-for-input transactions, set PROCLIM to a suitably high (or maximum) value so that you avoid rescheduling the application program. Use pseudo WFI and quick reschedule to eliminate processing overhead. By going through quick reschedule, a region can process more messages per schedule than the processing limit, thereby eliminating unnecessary rescheduling and reloading of application programs.

**Note:** A region scheduled against a transaction whose processing limit count is 0 cannot go through quick reschedule or become a pseudo WFI.

- Parallel scheduling in DB/DC and DCCTL environments

Do not use parallel scheduling for low-volume transactions that can be scheduled sequentially. Using message regions for parallel processing can reduce the efficiency of regions in which multiple transactions are processed for each scheduling of a program. Parallel scheduling should be reserved for those transactions where temporary peaks in the arrival rate might cause excessive queuing to develop due to “flooding” of a single message processing region. In these situations use a PARLIM level geared to the average service time and response time objective for the transaction type.

- IMS message processing regions in DB/DC and DCCTL environments

Operating your IMS system with too few or too many message processing regions can affect performance and transaction response times. In an environment with many different transaction types, regions with a very high occupancy level must be carefully monitored to ensure that they are not causing queues to build up on the IMS message queues.

In environments with only a few transaction types, high levels of occupancy can usually be sustained, because the application programs can often remain in the region processing multiple transactions for each schedule.

The meaning of high occupancy depends on the transaction mix, transaction characteristics, class scheduling scheme, and response-time objectives. For some transactions, a large queuing delay might be acceptable. For others, a much smaller queuing delay must be achieved in order to satisfy the user.

In contrast, having many low-occupancy concurrent message processing regions can unnecessarily increase the contention among regions for buffer pool space, real storage, IMS system data sets, and database records (because of program isolation enqueueing).

You can eliminate some of the processing overhead by specifying pseudo wait-for-input (PWFI) on the MPP region startup procedure. A PWFI message processing region remains scheduled when no more messages need to be processed, instead of terminating the application program. This eliminates unnecessary termination and rescheduling.

Adjust the number of message processing regions to the minimum possible consistent with achieving your response-time objectives. Increase the number of regions only when an increasing workload cannot be serviced within the objectives, and when changes to the message class structure, quick reschedule, PWFI, and region control parameters (including PARLIM, PROCLIM, and WFI) do not effectively contain the response-time increases. Although you can specify up to 999 dependent regions, always try to contain the response time demands with a minimum number.

The number of required regions, excluding wait-for-input regions, can be estimated as shown in the following example.

number of regions =  

$$((arrival\ rate) * (time\ in\ region)) / (desired\ region\ percentage\ occupancy)$$



*arrival rate* is the total number of transactions received in an interval divided by that interval in seconds. *time in region* is defined as the sum of the mean times for: scheduling and termination, schedule-to-first DL/I call, and elapsed execution, taken from the IMS Monitor Region Summary report.

Do not include regions dedicated to wait-for-input processing in this calculation, because their region occupancy is 100%.

In addition to establishing the number of message processing regions, you control which transactions can be scheduled into those regions by:

- The choice of class and priority structure
- The PROCLIM and PARLIM parameters on the TRANSACT macro
- CCTL regions in DB/DC and DBCTL environments

CCTL application program performance can be affected by contention for database resource adapter (DRA) resources. Specifically, if the maximum number of DRA threads allowed is less than the number of CCTL application programs making PSB schedule requests, some of those requests wait until a DRA thread becomes available.
- Save area sets in DB/DC and DCCTL environments

When you consider the number of regions and their occupancy for a given transaction rate, check that the input data transmission is not delayed. One factor causing delay might be the limiting number of concurrent terminal I/Os. You can adjust this limiting number using the SAV parameter in the online EXEC statement (up to a maximum of 999 concurrent I/Os). This parameter is termed the number of “dynamic save area sets” for the terminal I/O, because one of these areas controls each active terminal I/O.
- Database contention in DB/DC and DBCTL environments

Avoid databases with a small number of roots or root anchor points. If the databases are accessed by more than one transaction concurrently, program isolation (PI) contention is probable. The use of control records in databases or of records that keep running totals also causes PI enqueueing.

If you cannot avoid databases with a small number of roots, use class scheduling to allocate all such contending transactions to a single region, thus forcing them to process sequentially.

## IMS buffer pool tuning

You can configure the IMS buffer pools to improve overall system performance.

### Message queue pool tuning

The message queue pool acts as a buffer to the message queue data sets. Optimize the message queue pool to reduce unnecessary I/O activity.

If message queue I/O occurs during transaction processing, the message queue data sets can be detected as IWAIT items on the IMS Monitor Region IWAIT report. If the LRECL sizes of the short and long message queue buffers are not in correct proportion to each other, I/O can occur, because the mix of transactions uses up one or the other section of the buffer even though the size of the pool seems large enough.

The number of buffers is specified in the MSGQUEUE macro. You can override the number at execution time with the message queue buffer (QBUF) parameter in the EXEC statement. One tuning technique is to decrease the number of buffers until the monitor data shows increasing IWAIT instances marked as message queue I/O. Then you can use the QBUF parameter value plus one page-size unit.

You might want to identify unusual demands on message queue handling. Such demands can include:

- Large multi-segment input and output messages
- Extended processing between segment inserts
- Remote print output
- Terminal operator paging
- Multi-terminal output
- Excessive numbers of program-to-program switches
- Zero-priority messages for BMPs
- SPA sizes for conversational transactions

SPA sizes can adversely affect the performance of the message queues, detracting from the efficiency of single-segment messages that are used as input to well-designed application program processing.

The message queue LRECL size must allow for the SPA size, because the SPA is placed on the queue as the first segment of the message. The SPA segment can be split into multiple logical records when it is written to the IMS.LGMSG data set, or a larger LRECL size can be specified. This choice depends on the relative number of transactions requiring large SPA sizes.

You can use the Log Transaction Analysis utility to monitor the I/O queue times. The IMS Statistical Analysis utility produces a Line and Terminal report showing the average size sent and received, but does not show separate long and short message statistics. Average message lengths and counts for each transaction and terminal are part of the information in the system log type X'40' checkpoint records.

**Related reference:**

- ➡ Statistical Analysis utility (DFSISTS0) (System Utilities)
- ➡ Log Transaction Analysis utility (DFSILTA0) (System Utilities)
- ➡ DFSFIXnn member of the IMS PROCLIB data set (System Definition)

## **Message format pool tuning**

The message format buffer pool holds message format blocks for frequently used transactions. Optimizing it reduces the I/O burden of common transactions.

To obtain storage in the message format buffer pool, the space holding the least recently referenced format block is freed so that blocks for transactions with high arrival rates are usually available. To minimize format block I/O, plan to have sufficient space in the pool to hold the input and output format block pairs for use by your most frequently used transactions

Using fetch request elements (FREs) significantly affects efficient use of the pool. Each FRE controls one format block in the pool. If all the format blocks are 1000 bytes long and 10 FREs are specified, the maximum pool space is 10000 bytes. FRE assignments are made in the BUFPOOLS macro during system definition, but can be overridden by the FRE parameter in the control region EXEC statement. You can use the /DISPLAY POOL MFP command to compare the amount of free space to the amount of total space. If the free space is consistently high, relative to the size of the pool, make sure that the FRE assignment is not preventing full use of the allocated storage.



Another way to improve the use of the message format buffer pool is enable application programs to share formats. For example, a generalized message input format can share common DIF/MID blocks with subsequent editing performed by Input Message Field Edit and Input Message Segment Edit exit routines.

A performance-related function is available with Message Format Service (MFS). You can use the MFS Service utility to generate a directory of direct pointers to the location of all or a selected subset of the format blocks. This index is resident in the message format buffer pool and minimizes the I/O required to look up the location of a format block in the active IMS.FORMATA/B library directory. (Each index entry requires 14 bytes.) Evaluate the trade-off between allocating index storage for at least the highly referenced blocks against the reduction of I/O per block.

**Related reference:**

➡ MFS Service utility (DFSUTSA0) (System Utilities)

➡ DFSFIXnn member of the IMS PROCLIB data set (System Definition)

### **PSB, PSBW, and DMB pool tuning**

These guidelines for the program specification block (PSB) pool, PSB work (PSBW) pool, and data management block (DMB) pool can resolve pool space failures.

#### **PSB pool and intent list considerations**

The PSB pool holds all PSBs for scheduled application programs. A PSB remains in the pool until space is needed to load another PSB. Then, if all space is used, the least-referenced inactive PSB is freed. PSBs that are made resident and are not parallel scheduled are exceptions to this rule. These PSBs do not take from the pool allocation and are usually designated for highly referenced (usually preloaded) programs. You can page fix resident PSBs by including load list area DFSPSBRs in the module fix list coded for member DFSFIXxx in the IMS.PROCLIB data set.

If the DL/I address space option is used, two resident PSB pools exist, one in the z/OS common area (DFSPSBRs) and one in the DL/I address space (DFSIDLRS).

Similarly, resident-intent lists can be page fixed with the inclusion of DFSINTRS in the DFSFIXxx member. Make intent lists resident.

#### **Storage for the PSB work pool**

A portion of the PSB work pool is required by each active PSB, but the total size need not be larger than the maximum.

When you use the technique of starting large and reducing to some value greater than the maximum amount used, scheduling stops when any pool space failures occur.

#### **DMB pool considerations in DB/DC and DBCTL environments**

Make the DMB pool large enough to contain all IMS database DMBs. Each time you load a DMB into the pool, the database must be opened, and another database must be closed to free space. This activity increases the elapsed time that is required for DL/I calls. Limited page faulting is preferable to using OPEN and CLOSE. Make DMBs for active databases resident. These DMBs are good candidates to be page fixed. If all of the DMB pool is not being used, reduce the size until IWAIT for DMBs begin to occur. The IMS Monitor Region IWAIT report

shows the DMB pool size is too small by the presence of DMB= entries. You can then increase the specification of the DMB parameter on the BUFPOOLS macro.

### PSB pool, PSB work pool, and DMB pool space failures

Pool space failures occur when space in any pool is insufficient to sustain the current number of concurrently executing dependent regions. Each application program requires its PSB, the PSBW, the intent list, and the DMBs to be scheduled. Ensure that all pools are large enough to eliminate pool space failures. The general rule is to allow sufficient space for the  $(2R + 1)$  largest blocks that are part of the set DMB, PSB, or PSBW as appropriate, where R is the number of active dependent regions. This guidance allows for fragmentation in the pool space. Pool space requirements can be reduced by eliminating unnecessary PCBs from PSBs.

#### Related concepts:

 IMS internal resource usage (System Administration)

#### Related reference:

 DFSFIXnn member of the IMS PROCLIB data set (System Definition)

### Database buffer pool tuning

Allocate sufficient buffers in the IMS database buffer pools to prevent I/O that results from an application program having to reread data previously brought into the pool. Careful choice of subpool sizes and matching against database block sizes and frequency of database references minimizes unnecessary I/O.

You can either explicitly specify the Fast Path buffer pool size as part of the system definition process using the DBBF, DBFX, and BSIZ parameters, or you can use the Fast Path 64-bit buffer pool manager to control the number and size of the Fast Path database buffer subpools. To enable the Fast Path buffer pool manager, specify FPBP64=Y in the FASTPATH section of the DFSDFxxx member of the PROCLIB data set. When the Fast Path buffer pool manager is used, DEDB buffer pools are placed in 64-bit storage, while buffer pools for main storage databases (MSDBs), sequential dependent (SDEP) segments, system services, and buffer headers are managed in 31-bit storage.

Multiple OSAM subpools can be defined as having the same buffer size, and specific data sets can then be directed to specific subpools. If you have many OSAM data sets with similar or equal block sizes, you might be able to obtain some performance advantage by replacing a single large pool with separate subpools. Multiple smaller subpools reduce pool scanning to locate a segment. This method can also be used to prevent low-priority transactions (or BMPs) that access many database segments from monopolizing pool space.

**Tip:** The time required to sequentially process OSAM databases can be considerably reduced by using Sequential Buffering (SB).

You can define multiple VSAM local shared resource pools. If you have many VSAM data sets with similar or equal control interval sizes, you might get a performance advantage by replacing a single large subpool with separate subpools of identically sized buffers. Creating separate subpools of the same size for VSAM data sets offers benefits like those offered by OSAM multiple subpool support.

Use multiple local shared resource pools to specify multiple VSAM subpools of the same size. Create multiple shared resource pools and then place a VSAM subpool in each one that is the same size as other VSAM subpools in other local shared

resource pools. You can then assign a specific database data set to a specific subpool by assigning the data set to a shared resource pool. The data set is directed to a specific subpool within the assigned shared resource pool based on the control interval size of the data set.

You can also create separate subpools for VSAM KSDS index and data components within a VSAM local shared resource pool. Using these subpools can be advantageous, because index and data components do not need to share buffers or compete for buffers in the same subpool.

Multiple VSAM local shared resource pools enhance the benefits provided by Hiperspace buffering. Use Hiperspace buffering to allocate more buffers of 4 KB and multiples of 4 KB by using expanded storage in addition to virtual storage. Using multiple local shared resource pools and Hiperspace buffering allows data sets with certain reference patterns (for example, a primary index data set) to be isolated to a subpool backed by Hiperspace, which reduces the VSAM read I/O activity needed for database processing.

If you are using High Availability Large Databases (HALDBs), you can also direct HALDB partition data sets to specific OSAM or VSAM subpools.

**Related concepts:**

-  OSAM sequential buffering (Database Administration)
-  Hiperspace buffering (Database Administration)
-  Fast Path buffer performance considerations (Database Administration)

**Related reference:**

-  DFSVSMxx member of the IMS PROCLIB data set (System Definition)

## Trade-offs between I/O controlled by IMS and paging

When you are choosing between I/O controlled by IMS and paging, you should consider several trade-offs.

### General buffer pool considerations

Small tuned pools in IMS tend to minimize paging. In addition, because paging is not controlled by the IMS system, it can escalate dramatically in a storage-constrained environment.

In contrast, IMS buffer pool I/O is controlled and can be estimated and measured against those estimates. Finally, using a page fault and page I/O operation is generally more costly in processor terms than using a BSAM or IMS OSAM I/O.

Therefore, optimization of IMS in this area requires adjusting pool sizes while keeping in mind the costs in terms of paging and path length of scanning a pool.

### Page fixing the IMS buffer pools

Information for initial optimization of IMS buffer pools is explained in IMS buffer pool tuning (System Administration). Apply these guidelines to your system to achieve a minimum value for the total I/O operations (pool I/O and paging I/O). This means that you must consider a trade-off between I/O and storage across all pools, as well as within a single pool. For example, halving the maximum size of a large PSB pool might result in a very small increase in I/Os, but giving that extra storage to the message queue or format buffer pools might result in reducing even more I/O saving.

The way you tune your pools depends upon your real storage environment. In general, increasing maximum pool sizes increases paging; reducing pool sizes reduces paging of those areas.

If the IMS pools are subject to paging and IMS is sharing the processor usage, you can consider page fixing the significant buffers necessary for message and database access. This can be useful in a non-dedicated IMS system with very low-transaction rate, where IMS might be paged out unless the pools are page fixed. In heavily loaded IMS systems, response times might benefit from page fixing the data communication pools, because they hold data over the longest periods.

The principal IMS pools that can be page fixed are:

- Message queue pool (QBUF)
- DMB pool (DLDP)
- PSB pool (DLMP)
- Message format buffer pool (MFBP)
- Database subpools

If the DL/I address space option is used, two PSB pools exist: DLMP in the z/OS common area and DPSB in the DL/I address space.

Tell IMS to page fix all these areas, except the database subpools, by naming the resource in control statements that are included in the DFSFIXxx member of IMS.PROCLIB. The corresponding names used in these control statements are shown in parentheses in the above list.

To page fix the database subpools, you use control statements in the DFSVSMxx member of IMS.PROCLIB. The IOBF statement controls OSAM subpools. VSAM pool page fixing is specified with the VSAMFIX keyword of the OPTIONS control statement.

Usually, I/O for a pool impacts performance less than page faulting through the pool. If storage is a constraint, using a smaller fixed pool is preferable to using a larger pool that is subject to paging. However, page fixing is unlikely to improve performance, and might lead to higher overall paging rates, if your system is dedicated to IMS. This is because page fixing leaves a smaller available page pool for the rest of the system to use. Unless some other action is taken to reduce the storage requirements, higher rates inevitably result.

#### **Related reference:**

 DFSFIXnn member of the IMS PROCLIB data set (System Definition)

 DFSVSMxx member of the IMS PROCLIB data set (System Definition)

## **IMS log buffers**

When selecting the IMS online log data set (OLDS) block size and number of buffers in order to minimize system log I/O, similar concepts apply from general buffer pool considerations.

The log buffers only need to be plentiful enough to prevent IMS from waiting for a log write in order to get buffer space. The write-ahead data set (WADS) ensures that only full log buffers are written to the OLDS. The log write-ahead function does not truncate log buffers. When defining OLDS block sizes keep in mind that using a multiple of 4K will result in the log buffers being placed above the 2 gigabyte bar in real storage, but are below the 2 gigabyte bar in virtual storage.

Some buffer space is used for dynamic backout. When the backout records exist in the output buffer, they are used directly from the buffers. Otherwise, a buffer is removed from the set of output buffers and used to read the required log block from the OLDS. When backout is complete, the buffer is returned for use as an output buffer. Log buffers are long-term page fixed.

**Related reading:** See “General buffer pool considerations” on page 420 for more information on general buffer pool considerations.

### Using Library Lookaside under z/OS

As an alternative to IMS program loading with z/OS Program Fetch, you can specify the use of the z/OS service Library Lookaside (LLA). LLA is the recommended method for managing dynamically loaded program libraries.

For IMS application programs, it saves load modules in a data space, so the modules can be retrieved more efficiently than through DASD. LLA monitors the frequency with which modules are accessed and readjusts its data space population to optimize the probability of having a module in the data space when it is requested.

### Guidelines for application program preload in DB/DC and DCCTL environments

The preload option is commonly used for high-activity programs that do not use a large amount of virtual storage. You specify the names of the modules to be preloaded in member DFSMPLxx in the IMS.PROCLIB data set.

When the message processing region is initiated, you specify the suffix xx as a parameter in the EXEC statement of the DFSMPR procedure. Preloaded application programs are then branched to directly rather than through a FETCH program. Preloaded programs should have a schedule-to-first DL/I call elapsed time less than those that use the FETCH program, but page fault serialization can cause the application program elapsed time to increase.

**Restriction:** Do not preload application programs for batch regions executed using either the DLIBATCH or DBBBATCH procedures. There is no advantage to preloading application programs in these environments.

Guidelines for the most effective implementation of program preload are:

- All commonly used PL/I, VS Pascal, or COBOL subroutines and application program subroutines should be preloaded into each dependent region. If these subroutines are reentrant, they should be put into the pageable link pack area (PLPA) so that only one copy resides in real storage. Match subroutine usage with the region preload lists to ensure that the appropriate modules are preloaded.
- Application programs are candidates for preload when they account for a high percentage of a region transaction volume. Preload is most effective when the transaction arrival rate for the preloaded program is adequate to keep the working set of the program in real storage. If a system is constrained by real storage contention, preload only subroutines and very high-volume application programs because preloading can increase the paging rate.

In addition to deciding to preload them, consider class scheduling of high-volume transactions. Depending on the transaction arrival rate, it can be advantageous to preload in only one or two regions and class schedule the transactions accordingly.

The use of the preload option for system performance improvements is highly dependent on the availability of real storage and the arrival-rate of the candidate transactions.

- Application program overlay is sometimes a viable alternative to preload. If an all-purpose application program is coded to interrogate the input transaction data and execute only a portion of the application program code for the transaction subcode, program overlay I/O might be more efficient than loading or paging through the entire program. Preload is a better alternative if the transaction arrival rate is sufficient to maintain a working set in main storage.

## Minimizing path length

The total number of machine instructions that are executed to process a transaction, including system services, IMS services, and the application program itself, has a direct bearing on throughput. The accumulation of executed instructions is termed the *path length*. Options chosen for performance affect path length.

The actions suggested in “Choosing IMS options for performance” on page 412 all contribute to the minimization of path length. Avoid the regular use of traces such as the DL/I Call Image Capture and other traces invoked by the /TRACE command. These are specified as parameters on the OPTIONS statement in the DFSVSMxx member of IMS.PROCLIB.

Do not run the IMS Monitor (DFSMNTR0), except for during 10- to 20-minute preplanned intervals.

In a real storage-constrained system, the most effective way to reduce path length is to minimize paging. Minimal pools help to minimize paging by eliminating costly scanning of directories or buffers that might need to be paged in before they can be read. If virtual storage requirements are reduced:

- A minimal PSB pool minimizes buffer searching.
- A tuned database pool minimizes buffer searching; a larger database pool costs more in path length and might not reduce I/O.
- A tuned message queue pool minimizes buffer searching; a larger pool reduces IMS message queue I/O but does so at the expense of higher processor cycles per queue pool operation.
- The same applies to the message format buffer pool as to the message queue pool.

## Considerations for the communications network in DB/DC and DCCTL environments

Performance factors that affect response time can be related to NCP and network considerations or host transaction processing considerations.

- NCP and network considerations

Review the points in the network where delays might occur and try to avoid bottlenecks. Delays can occur:

- Within a cluster and queue for a line
- Waiting for transmission and waiting to be polled
- Through line errors
- Within the NCP
- Within VTAM



- Waiting for IMS to receive any macro
- Within IMS
- Waiting for transmission of output

Another possible problem is that VTAM cannot easily process larger output transmissions. You should check the appropriateness of the value for the OUTBUF keyword on the TERMINAL macro statements and ETO logon descriptor.

Input can also be delayed because IMS has no more input threads left.

- Host transaction processing considerations

Be aware of the trade-off made between the size of VTAM I/O buffers to service the communications network and the use of auxiliary storage. IMS communication traffic might share these resources with TSO and other subsystems.

Use the VTAM authorized path option to reduce processing of SEND and RECEIVE commands by bypassing validity checks of request parameter lists (RPLs) and I/O areas specified in the RPL. You do this by specifying VAUT=1 as a parameter on the EXEC statement for the control region.

IMS transactions should also be specified as response mode where possible. The major benefits are that the number of transmissions and line turnarounds is decreased and that line utilization is reduced.

## Guidelines for IMS system data set placement

In addition to tailoring the IMS system data sets to the workload, you should avoid undue contention for their use. Correct placement of high-usage data sets helps you avoid bottlenecks.

Here are some guidelines you should consider when placing your IMS online data sets:

- Separate among lightly used DASD volumes: heavily used libraries and data sets such as IMS.PGMLIB, the online log data sets, the active IMS.FORMATA/B and IMS.ACBLIBA/B libraries, and message queues.
- If high-activity data sets are on the same device, place them close to each other. The space allocation should stipulate contiguous storage and allocation by cylinder to avoid fragmentation of stored data.
- Have the volumes available to IMS mounted as private so that operating system temporary data sets are not allocated to these volumes.
- Do not place online IMS system data sets on shared DASD volumes. Avoid contention for both the device and the control unit. Position primary and secondary OLDSs on separate control units and channels when possible.
- To optimize program fetch I/O, order PGMLIB in descending frequency of use. Use full-track blocking. This minimizes Start I/O (SIO) operations and seek times.
- If using fixed-head devices, place write-ahead data sets (WADS) and database index data sets on the fixed-head portion.
- If program library contention is a problem, create multiple copies, each for a subset of the message regions. Do not mix production and test program libraries.
- For IMS shared data sets subject to updates, place the data sets on DASD devices capable of allowing parallel I/O from the same or multiple systems.

## I/O subsystem configuration

Use the RMF Channel Activity and Direct Access Device Activity reports to observe queuing delays in the I/O subsystem. Configure so that you minimize those delays by balancing the load across physical and logical channels.

## Application optimization in DB/DC and DCCTL environments

To reduce the amount of read I/O to the ACBLIB data set, consider loading the ACB members into 64-bit storage by specifying ACBIN64=ggg in the DATABASE section of the DFSDFxxx PROCLIB member. When the application is scheduled, the ACB members are retrieved from 64-bit storage instead of from the ACBLIB data set.

If possible, avoid the use of overlays in message processing programs. An exception might be if the application program uses large blocks of code exclusively for part of the transaction processing.

Do not have application programs that perform non-DL/I functions (such as STIMER and TTIMER, and contents supervision). Unless carefully controlled, such application programs usually perform poorly.

If lengthy application programs must be initiated online, consider splitting the application into two parts: an interactive part, which responds to the user and sends a program-to-program message to a batch part, which can be run in a low-priority region.

Consider redesigning applications whose performance is critical. Otherwise, plan to tune the system to process the transaction as it is before embarking on a total rewrite. Consider translating applications whose performance is critical from high-level language to assembler in order to improve program-load and execution time.

Always use the rollback (ROLB) call rather than cause the program to abnormally terminate. This keeps the message region available without operator intervention and with a smaller delay in message region availability to other transactions. (Abnormally terminating an application program keeps the region unavailable to other transactions waiting to be scheduled.)

## DL/I considerations

Consider the following things when designing DL/I databases.

- Use DL/I path calls wherever possible to reduce the DL/I path length and program execution time.
- Use qualified DL/I calls wherever possible to minimize the number of calls required to obtain the desired segment.
- Do not issue a GN call to the I/O PCB if the application program always retrieves single-segment input messages.
- Ensure that the application program returns to issue another GU to the I/O PCB on completion. This allows processing of another transaction of the same type if one exists on the input queue. This avoids the overhead of termination and rescheduling processing.
- When output is sent to several TP-PCBs, define multiple alternate PCBs rather than using the CHNG call.



- Select HDAM (or PHDAM, if using HALDBs) as the organization for a database wherever possible, because, when well tuned, HDAM (or PHDAM) generally results in shorter path lengths and fewer I/Os per DL/I call.

In assessing the characteristics of transaction processing and the interaction of application programs and databases, be aware of the performance design considerations presented in *IMS Version 11 Database Administration* and *IMS Version 11 Application Programming*.

---

## Planning for performance in a shared-queues environment

In a shared-queues environment, individual transaction performance might be poorer than in a non-shared-queues environment. However, because of increased parallelism and workload balancing, overall system performance and throughput should increase. And because input transactions are processed by the originating local IMS system (if a message region is available), you can influence IMS performance by controlling the number of dependent regions for each IMS system.

Each of the following can have an effect on system performance:

- IMS execution parameters: QBUF=, QBUFMAX=, LGMSGSZ=, and SHMSGSZ=
- CQS Local Structure Definition parameter, SYSCHKPT=
- CQS Global Structure Definition parameters: OBJAVGSZ=, OVFLWMAX=, and STRMIN=
- Size of the shared queues structures on the coupling facility
- Size of the z/OS system log structure on the coupling facility
- Number of z/OS system log data sets to back up the z/OS system log structure
- Frequency of structure checkpoints

The size of all log records has increased by eight bytes, and the size of Queue Manager and EMH log records has increased by an additional 32 bytes.

The SPA pool is no longer used, decreasing the amount of data logged for each checkpoint, and increasing the amount of available storage. The SPA is now maintained with its input and output messages.

For Fast Path systems, using the Fast Path Input Edit Router (DBFHAGU0) should be considered carefully because messages defined as “Local Only” will have priority over all other messages (“Local First” and “Global Only”).

---

## Identifying and correcting performance problems

Your general strategy is to decide the origin of the problem being experienced by performing a set of tasks.

Complete the following tasks to decide on the origin of the problem:

1. Isolate the problem into one of the following categories:
  - High paging rate seems to be slowing some or all of the workload.
  - High utilization of the processor by some subset of the total workload is causing dispatching problems for another (lower-priority) subset.
  - I/O contention in the DASD subsystem is slowing some subset of the workload.

- Utilization of, or contention for, resources associated with the communications subsystem is slowing the transmission of input or output messages.
  - Within IMS, utilization of physical resources (processor, I/O, storage) or contention for logical resources (pools, regions, control blocks, latches) has a negative impact on the performance of some or all of the transactions in one of the following areas:
    - Input or output message processing, including input queuing and Message Format Services (MFSs)
    - Program scheduling and termination
    - Program load and initialization
    - Program execution
2. Investigate further, if necessary, after you isolate the problem area, to determine:
    - The precise nature of the problem
    - The principal offenders by transaction or other category
    - The tuning action most likely to alleviate the problem
  3. Take the appropriate tuning action to prevent the problem from recurring.

## Examining paging rates

The operating system allocates real storage for the address spaces occupied by the control region or dependent regions. IMS might not obtain adequate machine cycles because of the excessive demand for real storage. Observe the paging rates and the frequency with which processing is delayed by a page fault. You can use the output from RMF II to examine the paging rates at peak IMS loads.

Page faulting in an IMS online system can directly increase the transaction time-in-system by the duration of the paging activity. If page faulting must occur, it is best to confine it to application programs in the dependent regions. A page fault experienced by the control region can hold up any dependent region waiting for a control region service. IMS provides some page-fixing options; these are covered in “Page fixing the IMS buffer pools” on page 420. However, if real storage is constrained, any page fixing generally serves to increase paging elsewhere and is not recommended. The primary tuning technique must be to reduce the virtual and the real storage requirement. This is particularly true in a dedicated IMS environment, where page fixing some portion of IMS only makes paging worse in the unfixed portion, possibly making overall paging rates worse.

You can monitor page rates dynamically, daily, or in detail. For more information about monitoring page rates, see “Deciding on monitoring activities and techniques” on page 375.

After the reason for the increased storage over commitment is identified, one or all of the following actions must be taken:

- Trade off IMS paging I/O against IMS controlled I/O.
- Reduce IMS paging I/O by reducing over commitment (reduce concurrent users) or by increasing IMS priority relative to other users.
- Speed up paging I/O by reducing I/O contention, splitting page data sets, or by placing page data sets on faster devices.
- In a non-dedicated IMS environment, try to reduce IMS paging, at the expense of other competing work, by fixing some portions of IMS.

## Detecting processor resource problems

In an environment with many jobs or subsystems operating concurrently, IMS competes for machine cycles. If those jobs or subsystems, such as an IMS batch execution or TSO, have a higher dispatching priority than the online IMS system, their workload has a marked effect on IMS performance.

**Note:** Monitor your resources dynamically, daily or in detail. For more information about how to perform this monitoring, see “Deciding on monitoring activities and techniques” on page 375.

## Tuning to remove I/O resource contention

A performance problem can occasionally be traced to I/O contention. Tuning activities for I/O devices, DASD storage access, channel usage, or control unit contention should not be confined to system-related devices. For IMS, an important factor is the volume of I/O activity, because reduced I/O activity reduces contention.

See “I/O subsystem configuration” on page 425 and the considerations for database I/O analysis included in “Program execution times” on page 434.

### Dynamic monitoring of the I/O subsystem

Typically, the I/O subsystem is not monitored dynamically, because corrective action normally requires reconfiguration of devices or data sets. This takes some time to plan and implement. However, if RMF is being run, the RMF Channel Activity and Direct Access Activity can be used to check against your objectives for I/O times, channel utilization, and device utilization.

Immediate action might not be possible. However, it might be possible to move data sets or reconfigure devices to correct the problem before the next IMS restart.

### Daily monitoring

Examination of the RMF Direct Access Device Activity and Channel Activity reports, together with knowledge of the I/O configuration and data set placement, might suggest a possible tuning action along the lines of those described in “Guidelines for IMS system data set placement” on page 424 and “Dynamic monitoring of the I/O subsystem.”

### Detailed monitoring

If additional data is required, plan to run GTF and obtain Volume Map, Seek Histogram, and System/Job Summary (EXCP-SIO-IO data) reports. These reports can be used to analyze I/O activity in great detail.

## Communication subsystem contention in DB/DC and DCCTL environments

Because of the great variety of transmission devices and types of lines, each suspected problem in the communication subsystem area must be separately examined.

As described in “Considerations for the communications network in DB/DC and DCCTL environments” on page 423, many possible causes of delay exist within the

network hardware and software. A variety of different tools might be required to investigate any suspected problems in this area.

## Dynamic monitoring

It is not possible to dynamically monitor the performance of the telecommunication subsystem.

## Daily monitoring

If the response time breakdown data indicates large and variable IMS output queue times, a bottleneck might exist in the IMS output message handling or on a line. Because it is not normally possible to reconfigure the TP network without considerable preparation time, plan to use the IMS Monitor and IMS line traces to investigate those IMS lines.

## Detailed monitoring

VTAMPARS can be used to analyze the GTF SMS records and to tune the VTAM buffer sizes. Sorting of DFSILTA0 and IMSPA Transit Log data by line or PTERM can isolate line-related problems for output queues to a single line or group of lines.

The IMS Monitor indicates line-related IWAITs, which require tuning of the IMS High Input/Output Pool (HIOP). Although HIOP is dynamically allocated, you might need to adjust either the buffer size definitions or the upper expansion limit. IMS Monitor output that is of interest for contention in communications are:

- The Communication Summary report gives the mean elapsed time for transactions using a VTAM node. This time can be affected by activity in the message format buffer pool.

The report also gives mean NON-IWAIT times. These can reflect device bottlenecks or problems in obtaining an input thread as described in Save area sets in DB/DC and DCCTL environments. The count of times that threads were not available is given in a separate report under the heading "REPORTS". The line 'Total Times ECBS waited for SAPS = nnn' gives the total. This figure is high if the monitor interval includes the time during which network initialization took place.

- The Line Functions report shows, for each active node, the data transmission activity, as described in Chapter 23, "Collecting and interpreting IMS monitoring data," on page 369. The TRANSMISSION BLKSIZE values show the number of characters passed by IMS to VTAM in OUTBUF buffers, a possible indicator of overloading of VTAM buffers. The TURNAROUND intervals can be used to identify inactive lines that might be eligible to distribute transaction entry.
- The Communication IWAIT report gives details of IWAITs for each node.

To trace VTAM activity, you can use GTF (including the Buffer Trace) to examine the content of VTAM buffers as data enters and leaves VTAM buffers. You can also use the Storage Management trace to reflect the overall use VTAM buffers.

Ultimately, tuning of network problems probably requires one or more of the following actions:

- Reduction of volumes of data transmitted. (See "Choosing IMS options for performance" on page 412.)
- Adjustment of VTAM or NCP options and parameters.

- Reconfiguration of the network to eliminate bottlenecks on multi-dropped lines.
- Isolation and correction of hardware delays caused by device errors.

## IMS message processing in DB/DC and DCCTL environments

This topic identifies possible performance problems associated with message processing and the message queues on an IMS system.

### Dynamic monitoring

MFS and Queue Pool statistics can indicate problems in this area. Pool sizes cannot be dynamically increased. They must be increased at the next IMS start. Normally, once this aspect of performance is tuned, it requires only periodic monitoring.

### Daily monitoring

If IMS output message processing is a problem area, IMS input message processing is probably affected by the same delays. Investigate any large or variable output queue time of IMS message queue and MFS (message format buffer pool) statistics using the IMS PA Resource Usage reports and DC Queue Manager Trace reports.

If output queue times are high in a DB/DC environment, check the IMS PA DC Queue Manager Trace report for output message lengths and run the IMS Monitor to study the Communications IWAIT report.

### Detailed monitoring

For detailed monitoring, use the following:

#### Message Queue Processing

IMS PA can be used to investigate the IMS Queue Manager in detail, using the DC Queue Manager Trace report. DFSUTR20 and IMSASAP give statistics on frequency and duration of communications-related IWAITs.

GTFPARS Job Summary and Detail Trace reports for the IMS control region can be used to examine the performance of MFS and Message Queue I/O.

#### Message Queue Data Set Tuning

The three data sets, IMS.QBLKS, IMS.SHMSG, and IMS.LGMSG, contain the directory and the short and long message segments for the message queues. In a busy online system with some constraints on the terminal I/O and conventional SPA pools, these data sets are the most active and critical. Place these data sets on different devices to minimize seek time. Check to see if the I/O activity is balanced between short messages and long messages. You might have to allow for a possible difference in device access times.

You can do this by examining the IMS Monitor Communication IWAIT report. The entries are organized by line number and reveal high frequency or high I/O elapsed times. Each line in the report identifies the DD name causing the IWAIT, so that you can pick out those that pertain to IMS.QBLKS, IMS.LGMSG, and IMS.SHMSG. The Region IWAIT report similarly identifies message queue I/O and totals IWAIT times for the region.

You can define up to 10 long and 10 short message queues.

**Related reading:** For more information on this option, see *IMS Version 11 System Definition*.

### **MFS Usage**

The GTFPARS Job Summary report for the IMS control region provides data on EXCP-SIO-IO timings to FORMATA/B. The IMS Monitor indicates in the Communication IWAIT report any IWAITs for Format Pool Space or for Directory or Block loading.

### **Message format library optimization**

The active IMS.FORMATA/B data set is the source of all MFS format blocks used during online execution and is usually a busy data set.

You can check the frequency or high I/O elapsed times against this data set using the IMS Monitor Communication IWAIT report. Report lines identified by I/O=DIR=fn or I/O=BLK=fn, where fn is the format block name, show I/O IWAIT time for directory lookup and block fetch by increasing the size of the format buffers. This pool size is specified with the FORMAT keyword on the BUFPOOLS macro for system definition. You can override the value for the online execution using the FBP parameter in the EXEC statement.

To reduce I/O to the directory, you can also create an index of track addresses of the blocks (or only those frequently referenced) and keep this index present in the MFS pool. The index is built when you execute the MFSRVC procedure, specifying INDEX on the utility control statement. You also include the names of the format or message blocks to be in the index.

Multiple reads of the active library data set occur when the size of the MFS control blocks stored on DASD is greater than either the track size or the z/OS block size.

## **Input queuing and scheduling/termination in DB/DC and DCCTL environments**

Another set of performance problems is brought about by failure to achieve the performance objectives. The problem might be observed by input queue buildup caused by either increased workload or application program scheduling delays.

### **Dynamic monitoring**

This topic explains how to obtain dynamic monitoring information for IMS input queuing, IMS region occupancy, IMS scheduling, and sync point processing.

#### **IMS input queuing**

Input queue time is not available. However, using the IMS /DISPLAY command with the QUEUE parameter, you can monitor the sizes of the input queues. You can also use the ACTIVE parameter to track performance and detect bottlenecks in scheduling and in the region occupancy area.

#### **IMS region occupancy**

Not available in any dynamic report. The IMS command /DISPLAY ACTIVE is used to monitor region usage.

#### **IMS scheduling**

The PSB and DMB Pools should already be tuned, as recommended in “Initializing z/OS and IMS parameters for tuning” on page 410. If all regions are busy, start any necessary additional regions. If some regions are idle, modify classes or PARLIM if one transaction is overloading a region.



**Related reading:** For more information, see “Avoiding contention for IMS resources (excluding buffer pools)” on page 414.

### **Sync point processing**

The occurrences of problems caused by waiting for synchronization points cannot be isolated dynamically.

## **Daily monitoring**

This topic explains how to obtain daily monitoring information for transaction response times, response time breakdown, region occupancy, program scheduling, and sync point processing at program termination.

### **Transaction response times**

Extract total (internal) response times for selected transactions that can be used as indicators of system performance.

**Related reading:** For more information on sources of response time data, see “Establishing performance objectives” on page 370.

### **Response time breakdown**

For similar transactions, extract response time breakdown (input queue, switch queue, processing, output queue) from IMS PA or DFSILTA0.

### **Region occupancy**

You can obtain indications of the relative use of the dependent regions in a DB/DC environment by using the IMS PA Resource Availability Reports. Also, you can directly use the summary of region occupancy percentages given in the IMS Monitor Region Summary report.

### **Program scheduling**

If the response-time breakdown data indicates large and variable input queue times, check the Region Occupancy figures. If all message regions have high occupancy, then another message region might be required. Alternatively, it might be possible to reduce occupancy by reducing program load or program execution times. If some or all of the IMS message regions are not busy, analysis of IMS PA Transaction Transit reports by transaction and class probably shows that one transaction or class is more critically hit than others. In this case, you should review the designation of classes and the allocation of classes to regions. PROCLIM and PARLIM should be reviewed also.

**Related reading:** For information on scheduling options, see “Choosing IMS options for performance” on page 412.

Programs executing as wait-for-input never show 100% occupancy even when they are in the region 100% of the time. Zero occupancy might be cause to review operator procedures, with instructions to manage the number of message regions based upon display of the queues.

The IMS PA Transaction Transit reports Graphic Summary is useful to analyze input queuing time by time of input across the entire measurement period. This can be used to discover if high input queue times result from a transient peak in transaction volumes or from a more sustained phenomenon. DFSILTA0 can be used for the same purpose, although its output is numeric rather than graphic.

### **Sync point processing at program termination**

If you are using IMS PA to provide response-time breakdown data, a high program switch queue time can indicate that delays are occurring because of a wait-for-synchronization point completion. This occurs if

MODE=MULT is used. Check the IMS PA CPU Usage report for multiple transactions of that type being dequeued per program schedule. As a general rule, use MODE=SNGL to avoid any delays of this type.

## Detailed monitoring

You should examine detailed monitoring information for scheduling and sync point processing.

### Scheduling

The IMS Monitor generates reports on PSB and DMB I/O counts and elapsed times. GTFPARS Job Summary and Detail Trace reports for the IMS control region can also be used for detailed investigation of these I/O events. However, scheduling problems are more often related to region availability and message class, PROCLIM, and region class assignments. These are described in “Avoiding contention for IMS resources (excluding buffer pools)” on page 414.

### Sync point processing

Delays caused by waiting for a synchronization point do not normally require detailed investigation.

## Program load and initialization

This topic describes dynamic and daily monitoring and detailed monitoring for program load and initialization.

### Dynamic and daily monitoring

The resource used to load or initialize a program cannot be isolated or corrected dynamically from normal daily monitoring data. Usually, this aspect of performance, once it is tuned, does not require more than periodic monitoring.

The IMS.PGMLIB data set is specified in the dependent region JCL and contains the application programs. The schedule-to-first DL/I call elapsed time shown on the IMS Monitor Region Summary and Program Summary reports includes the I/O time to bring a program into the region. Included in the I/O time is the system search through JOBLIB (searched first) or STEPLIB. If JOBLIB and STEPLIB are not used for the program library, IMS.PGMLIB should be first in the LNKSTnn member of SYS1.PARMLIB. Program libraries should always be blocked to full track size. Placement in the dependent region STEPLIB is recommended because a library can be built containing only those members required by the region.

In environments that are not constrained by real storage availability, consider making a trade-off between program loading and preload. You can use the schedule-to-first DL/I call elapsed times from the IMS Monitor Region Summary report to assess the potential savings of preload.

z/OS Library Lookaside (LLA) is the strategic method for managing dynamically loaded program libraries.

**Related reading:** For more information on LLA, see “Using Library Lookaside under z/OS” on page 422.



## Detailed monitoring

The GTFPARS Job Summary and Detail Trace reports for the IMS Message Regions can provide comprehensive data on the counts and sequences of SVC and I/O operations issued to locate and load the user's application program, and to execute the high-level language initialization modules.

For information about program library options and initialization considerations, see "Choosing IMS options for performance" on page 412.

## Program execution times

You can obtain program execution time information with dynamic monitoring of IMS internal response times, daily monitoring of program execution, and detailed monitoring.

### Dynamic monitoring of IMS internal response times

IMS internal response times are not available in any dynamic report. However, many installations have written special transactions that always perform a fixed (usually small) number of DL/I calls, such as a GHU, ISRT, REPL, DLET sequence. By use of this transaction, an IMS MTO can dynamically monitor the responsiveness of the IMS system. This is not a precise measure, but an indication of what a user is experiencing at that time.

### Daily monitoring of program execution

If the response-time breakdown data available from IMS PA or DFSILTA0 indicates large execution times, the cause might be related to one of the following:

- Program Load Time  
When daily monitoring indicates increases in execution times, the problem might be in the I/O subsystem.
- Program Initialization and Application Code  
After tuning, execution times are not likely to change significantly.
- DL/I Calls and I/O  
The IMS PA Internal Resource Usage report gives database I/O statistics. If these are high, the GTFPARS Data Set Overview report can be reviewed, together with the sizes and number of database buffers. Reorganize volatile databases regularly to reduce OSAM or VSAM overflow I/O. For further investigation of this area, the IMS Monitor should be run to obtain Call Summary and Region IWAIT reports.
- Other IMS Resource Delays  
The IMS Internal Resource Usage reports in IMS PA indicate if contention within program isolation activity is high, in which case, plans should be made to run and analyze the IMS PI trace.
- Dependent Region Modifications  
Examine the reasons for high elapsed time for dependent regions or for those application programs that exhibit excessive elapsed time for each transaction. You can use the IMS Monitor Program Summary report. Assuming that the dispatching priority is adequate, you can look for high schedule-to-first DL/I call elapsed times. Setting aside factors related to program load and paging, you need to find out if any modifications have been made in the dependent region code by your installation.

If changes have been made to support accessing z/OS data sets or the use of other z/OS services, you must evaluate the appropriateness and efficiency of these changes. One reason for these changes might be excessive overlay processing or initialization.

## Detailed monitoring

GTFPARS Job Summary and Detail Trace reports for the IMS Message Regions are the best method of examining the system-related activities of the application program in detail. IMSASAP Program Trace reports provide detailed analysis of application activity related to database access in a DB/DC environment. IMS PA Database Trace reports show breakdowns of database update activity in a DB/DC environment. DFSPIRP0 should be used to investigate any long delays if PI lockout is suspected as a possible cause.

The time spent in database retrieval that incurs I/O can be a large contributor to program execution time. The frequencies of each type of DL/I call to each database segment type and the frequencies of IMS IWAITS are given in the Call Summary report. If database-related delays are suspected as the cause of high program-execution times, that report helps identify the particular call and database combination that is responsible. The I/O wait times are shown in the IMS Monitor Region IWAIT report.

IWAIT times greater than 100 milliseconds are usually caused by interference from shared DASD or over commitment of resources. Databases using VSAM often experience shorter IWAITS, because the I/O buffer control managed by VSAM is able to meet the request without physical I/O. Excessive chaining of records to overflow tracks can be a contributing factor. Sometimes the data transfer rate of the storage device itself has a limiting effect.

In addition to database design considerations, the two principal strategies for decreasing I/O IWAIT time are:

- Reduce the frequency of I/Os by adding more buffers. This entails identifying inadequate IMS buffer allocation.
- Examine the I/O resource usage to discover if the distribution of data sets is a problem. This is a more protracted analysis that uses detailed output obtained from GTF and RMF traces.

A more extensive action is to change the DL/I structure or database organization.

**Related reading:** For more information on changing the DL/I structure or changing database organization, see *IMS Version 11 Database Administration*.



---

## Chapter 26. Making online changes

IMS supports two separate functions that support modifying IMS resources online: the online change function and dynamic resource definition.

- Online change function, of which there are two variations:
  - Local online change, which allows you to make online changes to IMS resources for one IMS.
  - Global online change, which allows you to coordinate online changes to IMS resources across all IMS systems in an IMSplex.
- Dynamic resource definition (DRD)

DRD enables you to make online changes to certain runtime resource definitions without using the online change process. To change other resources online (for example, resources in IMS.ACBLIB), you still need to use the online change process. With DRD, you can issue type-2 commands to create, update, and delete the runtime resource definitions dynamically, or use the enhanced Destination Creation exit routine (DFSINSX0), formerly called the Output Creation exit routine, to create transactions (and, if necessary, the programs that are associated with the transactions) instead of using the online change process.

You must ensure that:

- Commands are routed to and executed on the systems where you want the changes to be applied
- Any changes made dynamically with the DRD commands are recovered across a cold start
- The resource and descriptor definitions are exported to an RDDS

DRD enables you to create, update, query, and delete the following IMS resource definitions and their descriptor definitions dynamically, without using the system definition or online change processes:

- Application programs
- Databases
- Fast Path routing codes
- Transactions

Because DRD uses type-2 commands, you can only use DRD where your IMS is configured as follows:

- A single-IMS IMSplex that is configured with a minimal Common Service Layer that include an Operations Manager (OM) and a Structured Call Interface (SCI), but no Resource Manager (RM). This configuration is also known as an enhanced command environment.
- A multiple-IMS IMSplex that is configured with a full CSL (OM, SCI, and RM).

**Related reading:** For a detailed description of what you can do with DRD or information about configuring a single IMS with a minimal CSL, see *IMS Version 11 System Definition*.

The following topics describe the IMS online change function with DRD disabled.

### Related concepts:

 Overview of dynamic resource definition (System Definition)

### Related reference:

---

## Overview of the local online change function

In many installations, it is important that the online system be available during a large portion of the day. The ability to add, delete, and replace IMS databases, programs, transactions, and MFS formats online, without the necessity to bring down your IMS system, is a major step toward continuous operations. Adding, deleting, or changing IMS resources involves changes to the control blocks set up for these resources.

You can apply changes to the control block members within the IMS.FORMAT, IMS.ACBLIB, or IMS.MODBLKS data sets. You can apply changes to these data sets independently or in combination.

To use the IMS online change function, it is necessary to create three copies of each of the following libraries:

- IMS.MODBLKS—the library that contains the control blocks to support online change of databases, programs, transactions, and routing codes
- IMS.ACBLIB—the library that contains database and program descriptors
- IMS.FORMAT—the library that contains your MFS maps produced by the MFS Language and Service utilities

The libraries listed above are for the exclusive use of IMS offline functions and are called the staging libraries. For each library, a copy is made to produce a data set with a data set name suffixed with an A and a B, for example, IMS.FORMATA and IMS.FORMATB. These two copies of each library are used by the IMS online system.

At completion of IMS IVP processing the staging libraries and the IMS A libraries are identical. At this time, the A libraries are referred to as the active libraries. They are the libraries from which IMS draws its execution information. The B libraries are not used at this time and are referred to as the inactive libraries.

The following figure illustrates how libraries are used when you change your system online:

1. You apply changes to the staging libraries.
2. The staging libraries are subsequently copied to the inactive (B) libraries using the Online Change utility.
3. Operator commands are issued to:
  - a. Query the resources affected by the change to verify that the change is the correct one.
  - b. Cause the B libraries to become the active ones; the old active (A) libraries become the inactive ones.




Figure 46. How libraries are used when you change your system online

The process above is repeated as necessary.

You can change individual members of the ACB library instead of performing a full library switch by using a variation of the global online change process. For details of all variations of the global online change process, see “Overview of the global online change function” on page 444.

When you choose to add, replace, or delete any of the IMS resources mentioned in this topic, you apply your changes to the offline staging libraries by running one of the following:

- A MODBLKS type of system definition—if you have added, changed, or deleted applications, programs, full-function databases, DEDBs, routing codes, or transactions.

A MODBLKS system definition generates the control block members for resources that can be added or changed online. These control blocks are stored

in the library IMS.MODBLKS, and are used by the IMS control region, and the Multiple Systems Coupling Verification utility when an online change to your IMS system is requested.

- The ACB Maintenance utility (DFSUACB0), if you have added or changed any databases or programs.
- The MFS Language and Service utilities—if you have added or changed any MFS format definitions.

After the sequence of commands (/MODIFY for local online change or INITIATE OLC for global online change) has been issued to cause the previously inactive libraries to become the active libraries, your previously active libraries now become the inactive libraries. They are not destroyed until they are overwritten by the next online change sequence. You can return to the inactive libraries if backup and recovery are necessary, or if an incorrect definition occurs during your online change run.

Additionally, IMS monitors for you which set of libraries is currently active. If local online change is enabled, this information is kept in a status data set, IMS.MODSTAT. If global online change is enabled, this information is kept in the IMSPLEX.OLCSTAT data set.

After an online change is successfully completed, it persists across all types of IMS restarts. Additionally, the new resources can be easily maintained by running an SMP/E JCLIN against the Stage 1 output stream produced by your MODBLKS system definition to record the contents of the new system definition in your SMP/E control data set. This ensures that any maintenance applied to your IMS system is applied to the currently active IMS system. Do not manage the online change data sets with a migration/recall system that might recall the data set to a volume other than the one to which it was originally allocated. If you do so, IMS might be unable to warm start or emergency start the system.

## IMS environments that support the online change function

Global online change is not supported in all IMS environments that support local online change. Environments that require the MODSTAT/MODSTAT2 data sets do not support global online change. FDBR, XRF alternate and DBCTL-standby do not participate in global online process, however they keep track of the online change from the log records.

Use the following table to determine if global online change is supported in your environment.

*Table 34. IMS environments that support online change*

| Environment   | Supports local online change? | Supports global online change? |
|---------------|-------------------------------|--------------------------------|
| DB/DC         | Yes                           | Yes                            |
| DBCTL         | Yes                           | Yes                            |
| DBCTL-standby | No                            | No                             |
| DCCTL         | Yes                           | Yes                            |
| FDBR          | No                            | No                             |
| RSR active    | Yes                           | Yes                            |
| RSR tracker   | Yes                           | No                             |
| XRF active    | Yes                           | Yes                            |
| XRF alternate | No                            | No                             |

## Commands applicable to online change

To perform a local online change operation, use the /MODIFY PREPARE, /DISPLAY MODIFY, and /MODIFY COMMIT type-1 IMS commands.

You can also perform a local online change operation using type-2 IMS commands if you set up your IMS as a single-IMS IMSplex with a minimal Common Service Layer (CSL) consisting of a Structured Call Interface (SCI), an Operations Manager (OM), and no Resource Manager (RM). This IMS configuration is called the enhanced command environment. For information about defining your IMS with a minimal CSL (no RM), see *IMS Version 11 System Definition*. For information about performing a global online change operation for a single-IMS IMSplex (including performing an ACB library member online change), see “Overview of the global online change function” on page 444.

When you perform an online change, you make the inactive libraries (those containing the changes to be made) active, or copy members from the staging library to the active library. Use the following commands in the following sequence:

1. /MODIFY PREPARE

Use the /MODIFY PREPARE command to prepare an IMS for local online change and specify which resources to add, change, or delete. When you issue this command, IMS prevents terminals from queuing messages to database programs or transactions that will be changed or deleted.

You might have to issue additional commands after /MODIFY PREPARE to display and resolve work in progress that might cause the commit to fail.

2. /DISPLAY MODIFY

Use the /DISPLAY MODIFY command to produce a list of the resources (that are to be changed or deleted) that have work pending. No changes are accepted when work is pending. Use this command to ensure no work is pending before issuing a /MODIFY COMMIT command.

3. /MODIFY COMMIT

Use the /MODIFY COMMIT command to apply the changes specified in a preceding /MODIFY PREPARE command. When this command completes successfully, modifications persist across all IMS restarts.

If a /MODIFY PREPARE or a /MODIFY COMMIT command is unsuccessful and you do not want to retry the online change until another time, use the /MODIFY ABORT command to abort local online change on an IMS and reset the status to what it was before you issued the preceding /MODIFY PREPARE or /MODIFY COMMIT command.

## Changing the system definition online

You do not always need to rerun a complete IMS system definition to make changes to the IMS subsystem. Examine the changes to see if an online change can be made.

If the request does not involve changes to the IMS network or to static terminals, you can probably make the change online.

The procedure to make changes to the IMS online system can be summarized as:

1. If necessary, make required system definition Stage 1 input changes.
2. If necessary, perform a MODBLKS system definition.



3. If necessary, run the MFS Language and MFS Service utilities.
4. If necessary, generate DBDs, PSBs, and ACBs.
5. As necessary, use the Online Change Copy utility to produce suitable copies of the IMS.MODBLKS, IMS.ACBLIB, and IMS.FORMAT, data sets.
6. Issue the appropriate online change command sequence to make the online change at the proper time. The changes must be copied from the inactive library for a library switch OLC.
7. Verify that the required changes took place.

If you need to undo or back out an online change, repeat the online change command sequence without changing the inactive libraries.

## Changing programs and transactions

In an IMS DB/DC or DCCTL environment, the online change PREPARE command stops all incoming messages and transactions for a resource to be changed, but does not affect messages that have already been queued. The online change COMMIT command does not take effect until messages for a directly affected resource to be changed or deleted have been scheduled (dequeued), so you must either wait for the messages to be scheduled or stop the associated transactions.

If, for valid reasons, you do not want to wait for the messages to be scheduled (for the online change COMMIT command to take effect), use the /STOP or /PSTOP command to stop the associated transactions before issuing the online change COMMIT command.

**Recommendation:** Plan to make your changes when activity associated with the resource is low—when there are few transactions to be processed. When activity is low, you will not have to wait long for messages to be scheduled or have many transactions to stop if you do not want to wait.

You must stop all message-driven Fast Path programs and wait-for-input programs that directly refer to resources to be changed or deleted because the online change COMMIT command does not execute if any wait-for-input programs have been scheduled.

Also, any active conversations involving resources to be changed or deleted must complete before the online change COMMIT command can complete.

## Implementing security changes online

While IMS is active, changes to RACF security definitions are made completely outside of IMS.

The following general procedure describes how to make changes to RACF:

1. In the RACF database, update the RACF security profiles that describe IMS resources (such as transactions and commands) by issuing RACF commands.
2. Issue the RACF command SETROPTS RACLIST(classname) REFRESH to refresh the profiles for the appropriate security class in the RACF data space.

## Changing MFS formats

You can use the online change COMMIT command to change MFS formats without stopping related activity (messages can be queued and scheduled while you change an associated MFS format block).

But you should coordinate changing MFS formats and message queuing so that problems do not occur.

In the IMS DB/DC and DCCTL environments, the online change PREPARE command does not stop the queuing of related messages when you change MFS formats.

## Making online changes for DEDBs

In the IMS DB/DC and DBCTL environments, you can make both database-level and area-level changes to Fast Path data entry databases (DEDBs).

A database-level change affects the structure of the DEDB, and includes such changes as adding or deleting an area, adding a segment type, or changing the randomizer routines. An area-level change involves increasing or decreasing the size of an area, including the overflow portions (IOVF and DOVF) and the control interval (CI) size.

A database-level change requires you to stop all areas of the DEDB. An area-level change requires you to stop the area using either the /DBRECOVERY AREA command.

IFP and MPP regions that access the changed DEDBs do not need to be stopped during the online change.

**Related reading:** For more information about adding, modifying, or deleting DEDBs online, see *IMS Version 11 Database Administration*.

## Making online changes for HALDBs

The following restrictions apply for IMS online change with HALDB partitions.

- During an ALL or MODBLKS system definition, you cannot redefine a database that is currently defined as a HALDB partition without performing a cold start of IMS. A cold start is required even if you delete the database online. That is, if you delete a HALDB partition, you cannot reuse that database name for any database that you plan to add during an ALL or MODBLKS online change without an IMS cold start.
- You cannot convert a database that is defined during an ALL or MODBLKS system definition into a HALDB partition without performing a cold start of IMS. A cold start is required even if you delete the database online. After the IMS cold start, you must redefine the database as a HALDB partition.

These restrictions do not apply for a HALDB master database.

## Making local online changes in an XRF complex

In the IMS DB/DC and DCCTL XRF environments, in addition to preparing the staging libraries, you need to prepare the inactive libraries.

If an XRF takeover occurs while the active IMS is processing a /MODIFY COMMIT command, the change might or might not have occurred. At the end of an online change, IMS writes a log record; the XRF alternate subsystem reads the log and makes the online changes to match the changes on the active subsystem. Use the /DISPLAY DB or QUERY DB command to check that a change occurred on the alternate subsystem.

## Online change performance considerations

Deleting databases from the IMS.MODBLKS data set using Online Change incurs the performance overhead of reading the intent lists from the ACB library for all application programs in MODBLKS whose intent lists are not loaded.

This is necessary to determine which transactions need to stop queuing, to prevent work that might cause commit to fail. This performance overhead is proportional to the number of application programs defined in MODBLKS for which the intent list is not loaded.

**Recommendation:** If you need to delete databases from MODBLKS, consider performing an IMS cold start instead of using local or global online change.

---

## Overview of the global online change function

Global online change is an IMS function to change resources for all IMS systems in an IMSplex. A Master IMS control region, as specified by the user or by an Operations Manager (OM), uses Resource Manager (RM) to coordinate the phases of online change with other IMS systems in an IMSplex.

For member online change, the member names specified on INIT OLC TYPE(ACBMBR) NAME(xxx) and their associated members are copied from staging to active. Staging may have members that are not copied. You can add or change individual members of IMS.ACBLIB and bring these members online without quiescing the IMSplex or requiring a full refresh of the active ACB library.

For library switch online change (local or global) the active and inactive libraries are swapped.

For more information about changing or adding members to IMS.ACBLIB online, see “Changing or adding IMS.ACBLIB members online” on page 454.

If you are running an IMSplex that does not use RM, you can specify the CSL global parameter RMENV=N in the DFSCGxxx IMS.PROCLIB member. If this parameter is specified, IMS does not require RM and will not use RM services, but type-2 command support is available. If OLC=GLOBAL, use the INITIATE OLC command to initiate local online change. The OLCSTAT data set must be defined. This data set cannot be shared between IMS systems.

With global online change, the Master IMS control region coordinates online change of the following resources among all the IMS systems in an IMSplex:

- Databases (DMB in IMS.ACBLIB)
- Database directories (DDIR in IMS.MODBLKS)
- MFS formats (IMS.FMTLIB)
- Programs (PSB-type ACBs in IMS.ACBLIB)
- Program directories (PDIR in IMS.MODBLKS)
- Transactions (SMBs in IMS.MODBLKS)

## Commands applicable to global online change

To perform a global online change, issue a command to make the inactive libraries (those containing the changes to be made) active or to copy members from the staging library to the active.

The following commands make the inactive libraries active or copy the staging library to the active:

1. **/DISPLAY MODIFY**

This command produces a list of the resources (that are to be changed or deleted) that have work pending. No changes are accepted when work is pending. Use this command to ensure no work is pending after issuing an **INITIATE OLC PHASE(PREPARE)** command and before issuing an **INITIATE OLC PHASE(COMMIT)** command.

2. **INITIATE OLC PHASE(PREPARE)**

This command prepares all of the IMS systems in the IMSplex for global online change by specifying which resources to add, change, or delete.

When running in an environment without RM (RMENV=N), this command performs the online change only at the command master IMS or the single IMS to which the command is routed.

3. **INITIATE OLC PHASE(COMMIT)**

This command commits global online change on all IMS systems in an IMSplex.

When running in an environment without RM (RMENV=N), this command performs a local online change on the IMS that processes it.

You might have to issue additional commands after **INITIATE OLC PHASE(PREPARE)** to display and resolve work in progress that might cause the commit to fail. These commands include:

**TERMINATE OLC**

This command aborts global online change on all IMS systems in an IMSplex.

When running in an environment without RM (RMENV=N), this command terminates a local online change on the IMS that processes it.

**QUERY MEMBER**

This command displays status or attribute information about IMS systems in an IMSplex, including local or global online change status, if applicable.

**QUERY OLC**

This command provides information about the **OLCSTAT** data set contents or resources being modified with **TYPE(ACBMBR) OLC**.

## Changing the system definition online

You do not always need to rerun a complete IMS system definition to make changes to the IMS subsystem. Examine the changes to see if an online change can be made.

If the request does not involve changes to the IMS network or to static terminals, you can probably make the change online.

To make changes to the IMS online system:

1. If necessary, make required system definition Stage 1 input changes.
2. If necessary, perform a **MODBLKS** system definition.
3. If necessary, run the MFS Language and MFS Service utilities.
4. If necessary, perform **DBDGEN**, **PSBGEN**, and **ACBGEN**.
5. As necessary, use the Online Change Copy utility to produce suitable copies of the **IMS.MODBLKS**, **IMS.ACBLIB**, and **IMS.FORMAT** data sets.

6. Issue the appropriate online change command sequence to make the online change. The changes must be copied from the inactive library for a library switch OLC.
7. Verify that the required changes took place.

If you need to undo or back out an online change, repeat the online change command sequence without changing the inactive libraries.

## Changing programs and transactions

In an IMS DB/DC or DCCTL environment, the online change PREPARE command stops all incoming messages and transactions for a resource to be changed, but does not affect messages that have already been queued.

The online change COMMIT command does not take effect until messages for a directly affected resource to be changed or deleted have been scheduled (dequeued), so you must either wait for the messages to be scheduled or stop the associated transactions.

If you do not want to wait for the messages to be scheduled (for the online change COMMIT command to take effect), use the /STOP or /PSTOP command to stop the associated transactions before issuing the online change COMMIT command.

**Recommendation:** Plan to make your changes when activity associated with the resource is low—when there are few transactions to be processed. When activity is low, you will not have to wait long for messages to be scheduled or have many transactions to stop if you do not want to wait.

You must stop all message-driven Fast Path programs and wait-for-input programs that refer to resources to be changed or deleted because the online change COMMIT command does not execute if any wait-for-input programs have been scheduled.

Also, any active conversations involving resources to be changed or deleted must complete before the online change COMMIT command can complete.

## Implementing security changes online

While IMS is active, you must make any changes to RACF security definitions completely outside of IMS.

The following general procedure describes how to make changes to RACF:

1. In the RACF database, update the RACF security profiles that describe IMS resources (such as transactions and commands) by issuing RACF commands.
2. Issue the RACF command SETROPTS RACLIST(classname) REFRESH to refresh the profiles for the appropriate security class in the RACF data space.

## Changing MFS formats

In the IMS DB/DC and DCCTL environments, the online change PREPARE command does not stop the queuing of related messages when you change MFS formats.

You can use the online change COMMIT command to change MFS formats without stopping related activity (messages can be queued and scheduled while you change an associated MFS format block), but you should coordinate changing MFS formats and message queuing so that problems do not occur.

## Making online changes for DEDBs

In the IMS DB/DC and DBCTL environments, you can make both database-level and area-level changes to Fast Path data entry databases (DEDBs). A database-level change affects the structure of the DEDB, and includes such changes as adding or deleting an area, adding a segment type, or changing the randomizer routines.

An area-level change involves increasing or decreasing the size of an area, including the overflow portions (IOVF and DOVF) and the control interval (CI) size.

A database-level change requires you to stop all areas of the DEDB. An area-level change requires you to stop the area using the UPDATE AREA STOP(ACCESS) command.

IFP and MPP regions that access the changed DEDBs do not need to be stopped during the online change.

**Related reading:** For more information about adding, modifying, or deleting DEDBs online, see *IMS Version 11 Database Administration*.

## Making online changes for HALDBs

The following restrictions apply for IMS online change with HALDB partitions.

- During an ALL or MODBLKS system definition, you cannot redefine a database that is currently defined as a HALDB partition without performing a cold start of IMS. A cold start is required even if you delete the database online. That is, if you delete a HALDB partition, you cannot reuse that database name for any database that you plan to add during an ALL or MODBLKS online change without an IMS cold start.
- You cannot convert a database that is defined during an ALL or MODBLKS system definition into a HALDB partition without performing a cold start of IMS. A cold start is required even if you delete the database online. After the IMS cold start, you must redefine the database as a HALDB partition.

These restrictions do not apply for a HALDB master database.

## Changing runtime resource definitions

You can create, update, and delete IMS resources and their descriptors dynamically by using DRD. However, you must use the ACB member online change function to change ACB members.

ACB members can be associated to a database or program resource as part of a CREATE DB or CREATE PGM command or with an online change. There are two ways that you can associate ACB members in an IMSplex:

- Create all ACBs in an ACBLIB and then use DRD to associate database and program resources to the ACB.
- Create the database and program resources using DRD commands and then perform an online change to associate the resources to the ACB.

CREATE commands for database and program resources do not require that complimentary ACB members exist in the ACBLIB. If ACB members exist for resources that are being created, the association is created between the resource and the ACB member when the CREATE command is issued, and an online change is not required.



### **Associating an ACB member using DRD commands**

After running ACB generation, use the CREATE DB or CREATE PGM command to create a new database or program resource.

To associate an ACB member as part of the DRD CREATE command, complete the following steps:

1. Run ACB generation to create an executable ACBLIB member in a staging ACB library.
2. Issue the INITIATE OLC command to bring the ACB member into the active ACBLIB.
3. Issue the CREATE DB or CREATE PGM command to create a new database or program resource. The database is automatically associated to the ACBLIB through the CREATE command.

### **Associating an ACB member using the ACB member online change function**

After running ACB generation, issue the INITIATE OLC command to associate the ACB member to the database or program resource.

To associate an ACB member with the ACB member online change function, complete the following steps:

1. Issue the CREATE DB command to create a new database.
2. Run ACB generation to create an executable ACBLIB member in a staging ACB library.
3. Issue the INITIATE OLC command with the TYPE(ACBMBR) parameter to copy the updated or new member from the staging library to the active library and associate the ACB member to the database or program resource.

### **Associating an ACB member using the ACB online change function**

After running ACB generation, issue the INITIATE OLC command to associate the ACB member to the database or program resource.

To associate an ACB member with the ACB online change function, complete the following steps:

1. Issue the CREATE DB command to create a new database.
2. Run ACB generation to create an executable ACBLIB member in a staging ACB library.
3. Issue the INITIATE OLC command to perform a library switch online change and associate the ACB member to the database or program resource.

## **Making local online changes in an XRF complex**

In the IMS DB/DC and DCCTL XRF environments, in addition to preparing the staging libraries, you must prepare the inactive libraries.

If an XRF takeover occurs while the active IMS is processing a INIT OLC PHASE(COMMIT) command, the change might or might not have occurred. At the end of an online change, IMS writes a log record; the XRF alternate subsystem reads the log and makes the online changes to match the changes on the active subsystem. Use the /DISPLAY DB or QUERY DB command to check that a change occurred on the alternate subsystem.

## Making global online changes in a sysplex that uses XRF

The XRF alternate does not participate directly in a global online change. The alternate participates indirectly once it receives the X'70' and X'22' log record written by the XRF active system.

The QUERY MEMBER command might show some online change status for the XRF alternate while it is processing the X'70' and X'22' log record.

During takeover of an active system, the XRF alternate updates the OLCSTAT data set with its IMSID. After the takeover is complete, the XRF alternate behaves like any IMS active system and can participate in a global online change. The abended XRF active system can be brought back up as the new XRF alternate, if desired.

The MODSTAT and MODSTAT2 data sets are not used. You can leave the MODSTAT and MODSTAT2 data set DD statements in your IMS control region JCL, but they are ignored.

Global online change in an XRF environment has the following limitations:

- If an XRF active system fails or is shut down (/SWITCH SYSTEM ACTIVE or /CHE FREEZE) and the XRF alternate system takes over, IMS deletes the IMSID of the failed or shut down XRF active system from the OLCSTAT data set during takeover processing, so that global online change can work even if the old active is down. After takeover, if you bring up the failed or shut down XRF active system as the new XRF alternate system, and initiate global online change, the new XRF alternate system does not process the online change command directly, but performs the online change later when it detects the online change log record.
- If the XRF active system fails and the XRF alternate system attempts to take over before it catches up with the XRF active system's global online change state, the XRF alternate system abends with U2801 RC=X'0011'. The XRF alternate is not synchronized with the OLCSTAT data set. Because other IMS systems in the IMSplex might be synchronized with the OLCSTAT data set, the XRF alternate is not permitted to takeover.

## Global online change after XRF takeover and DBCTL standby emergency restart

During the XRF takeover, the IMSID of the old active is removed from the OLCSTAT and the IMSID of the new active is added to the OLCSTAT. If there was any error removing the active's IMSID, the takeover completes, but two IMSIDs will exist in the OLCSTAT (the old XRF active and the new XRF active IMSIDs).

After an emergency restart of a DBCTL standby with RMENV=N and OLC=GLOBAL, two IMSIDs exist in the OLCSTAT data set (the old DBCTL active and new DBCTL active IMSIDs).

You can run the Global Online Change utility (DFSUOLC0) with FUNC=DEL to delete the old active XRF or old DBCTL active IMS systems IMSID from the OLCSTAT data set. Otherwise, the next INITIATE OLC PHASE(PREPARE) command you issue to initiate global online change will fail if OPTION(FRCABND,FRCNRML) is not specified.

**Related reading:** For more information about XRF and DBCTL standby, see "Overview of XRF" on page 585.



## Online change performance considerations

Depending on what you are using the online change function for, different performance considerations apply.

### Deleting databases from the MODBLKS data set

Deleting databases from MODBLKS using Online Change incurs the performance overhead of reading the intent lists from ACBLIB for all application programs in MODBLKS whose intent lists are not loaded.

This is necessary to determine which transactions need to stop queuing, to prevent work that might cause commit to fail. This performance overhead is proportional to the number of application programs defined in MODBLKS for which the intent list is not loaded.

**Recommendation:** If you need to delete databases from MODBLKS, consider performing an IMS cold start instead of using local or global online change.

**Related reference:**

 INITIATE OLC command (Commands)

## Making local online changes in a sysplex

IMS does not automatically coordinate local online changes across a sysplex. Use global online change to coordinate online change across a sysplex. This topic describes how to manually coordinate local online changes across a sysplex.

**Recommendation:** Clone your IMS subsystems and share libraries (ACBLIB, MODBLKS, FORMAT) across the sysplex. MODSTAT IDs must be the same across the sysplex.

To prepare for performing IMS online change in a sysplex environment:

1. Use an Automated operator (AO) application program. The AO application program should be able to:
  - Stop resources you want to change or delete.
  - Handle online change exception conditions for an individual IMS subsystem.
  - Issue the /MODIFY PREPARE or /MODIFY ABORT commands.

Design the AO exit routine to look for the online change messages that IMS issues and to take the suggested action. Design an AO application program or use an automation tool to issue commands to resolve any online change problems.

2. Define the active, inactive, and staging libraries.
3. Share the active and inactive libraries.

All IMS subsystems that you want to be able to change online should also share the staging libraries.

4. Define each IMS (with its own MODSTAT data set).

Each IMS must have its own MODSTAT data set, with the same active and inactive ACBLIB, FMTLIB, and MODBLKS definitions.

5. Start each IMS.

To perform the online changes:

1. Perform system definitions on the staging libraries for changes.
2. Run the Online Change Copy utility.

3. Issue the /MODIFY PREPARE command for each IMS subsystem.  
You can have either the IMS operator or an AO application program issue this command. If the command fails on any IMS subsystem in the sysplex, the operator or AO application program must issue the /MODIFY ABORT command on all IMS subsystems for which the /MODIFY PREPARE succeeded.
4. **Optional:** Stop all resources that will be changed or deleted.  
This step greatly increases the chance that the online change will succeed for an individual IMS subsystem.
5. Resolve any work in progress.
6. Issue the /DISPLAY MODIFY command.  
Use the /DISPLAY MODIFY command to display the work in progress for resources to be changed or deleted, before attempting the COMMIT phase. When COMMIT is successful, the modifications persist across all IMS restarts, unless global online change occurs while an IMS is down.
7. Issue /MODIFY ABORT|COMMIT for each IMS subsystem.
8. Start resources that were changed or added.

## Enabling the global online change function

Use the Global Online Change utility (DFSUOLC) to initialize the OLCSTAT data set and enable an IMSplex for global online change.

To enable an IMSplex for global online change, you must:

1. Run the Global Online Change utility (DFSUOLC0) to initialize the OLCSTAT data set.

The Global Online Change utility needs to be run once before the first IMS cold starts the first time. The OLCSTAT data set is comparable to the MODSTAT data set used by local online change.

### Important:

- Use the Global Online Change utility with caution, so you do not inadvertently destroy valid OLCSTAT data set contents. If you do destroy valid OLCSTAT data set contents, global online change and initialization of additional IMS systems fail until the OLCSTAT data set is re-initialized.  
You should establish an OLCSTAT data set recovery procedure to deal with the loss of the OLCSTAT data set. After every successful global online change, record the modify ID, the active online change library suffixes, and the list of IMS systems that are current with the online change libraries. If the OLCSTAT data set is destroyed, run the initialize function of the Global Online Change utility with the saved data to re-initialize the OLCSTAT data set.

**Related reading:** For more information about the Global Online Change utility, see *IMS Version 11 System Utilities*.

- If you build a multisystem sysplex, you automatically build a global resource serialization complex with it. Global online change relies on global resource serialization to serialize access to the OLCSTAT data set from multiple systems. If you use global online change in a sysplex with more than one z/OS system, you should not define the GRSRNL=EXCLUDE parameter in the MVS parmlib. Defining GRSRNL=EXCLUDE makes the OLCSTAT data set serialization local instead of global.
- For an IMSplex running without an RM (with CSL global parameters RMENV=N and OLC=GLOBAL), you must ensure that another IMS

- subsystem does not attempt to use the OLCSTAT data set from the IMS subsystem that is running without RM.
2. For each IMS in the IMSplex, perform these steps:
    - a. Remove the MODSTAT DD and MODSTAT2 DD statements from the IMS control region JCL.
    - b. Define DFSCGxxx IMS.PROCLIB member parameters related to global online change, or in the CG section of DFSDFxxx member.  
Specify OLC=GLOBAL to enable global online change.  
Specify OLCSTAT= with the online change status data set name. All of the IMS systems in the IMSplex must specify the same OLCSTAT data set.
    - c. Shut down IMS.
    - d. Cold start the IMS.

## Initiating a global online change

The INITIATE OLC (initiate online change) command is provided to support global online change, where online change is coordinated across IMS systems in the IMSplex. INITIATE OLC is similar to /MODIFY PREPARE and /MODIFY COMMIT, except that it applies to an IMSplex-wide global online change. OM sends the INITIATE OLC command to one IMS in the IMSplex.

For syntax, environment, and usage information, see *IMS Version 11 Commands, Volume 1: IMS Commands A-M*.

**Note:** When running in an environment without RM (RMENV=N), this command performs a local online change on the IMS that processes it.

INITIATE OLC prepares for or commits a global online change of IMS resources across an IMSplex.

INITIATE OLC is not supported if local online change is enabled. The INITIATE OLC PHASE(PREPARE) command is rejected if the IMS to which the command is routed does not support global online change. If the INITIATE OLC PHASE(PREPARE) command is rejected and there is an IMS that supports global online change, route the command to an IMS that supports global online change.

The correct online change command sequence is the following:

1. INITIATE OLC PHASE(PREPARE)
2. /DISPLAY MODIFY Use the /DISPLAY MODIFY command to display the work in progress for resources to be changed or deleted, before attempting the COMMIT phase. When COMMIT is successful, the modifications persist across all IMS restarts, unless global online change occurs while an IMS is down.
3. INITIATE OLC PHASE(COMMIT)

If the INITIATE OLC PHASE(PREPARE) is specified without a FRCABND or FRCCNRML keyword and the command fails because one or more IMS are down or go down before the online change is committed, the online change must be aborted and started over. Issue the TERMINATE OLC command to abort the online change.

If the INITIATE OLC PHASE(PREPARE) is specified with the FRCNRML keyword and the command fails for any IMS, you can proceed with an INITIATE OLC PHASE(COMMIT) command after shutting down those IMS systems where the prepare failed. Otherwise you must abort the online change and start over.

If the INITIATE OLC PHASE(PREPARE) is specified with the FRCABND keyword and the command fails for any IMS, you can proceed with an INITIATE OLC PHASE(COMMIT) command after canceling those IMS systems where the prepare failed. Otherwise, abort the online change and start over.

If the INITIATE OLC PHASE(COMMIT) command fails for any IMS before the OLCSTAT data set is updated, you can either correct the errors and try the commit again or abort the online change with a TERMINATE OLC command.

If the INITIATE OLC PHASE(COMMIT) command fails for any IMS after the OLCSTAT data set has been updated, you can correct the errors and try the commit again. The online change cannot be aborted.

If an IMS abends during online change and the INITIATE OLC PHASE(PREPARE) command was not specified with FRCABND, then issue the TERMINATE OLC command to abort the online change. The INITIATE OLC PHASE(COMMIT) command is not permitted in this case. If an IMS abends during online change and the INITIATE OLC PHASE(PREPARE) command was specified with FRCABND, then the INITIATE OLC PHASE(COMMIT) command is permitted.

Type-1 commands and type-2 commands that come from the OM interface are rejected during the commit phase, if the command changes resources. Commands that change resources could interfere with the online change of the resources. Type-1 commands and type-2 commands that come from the OM interface are permitted during the commit phase, if the command (such as /DISPLAY or QUERY) displays resources. Type-1 commands that are entered from the system console or an IMS terminal are queued during the online change commit phase. These commands run after the online change is committed or aborted.

You can specify INITIATE only through the OM API.

INITIATE is invalid on an XRF alternate, RSR tracker, and FDR system.

Each IMS participating in the global online change does not issue the same synchronous online change messages to the master terminal or system console that it does for a local online change. The OM command response contains information equivalent to the online change messages that appear for the local online change, such as the DFS3499 message contents.

Each IMS participating in the global online change may issue asynchronous online change messages to the system console, such as DFS3400, DFS3445, and DFS3498. For more information about these messages, see *IMS Version 11 Messages and Codes, Volume 1: DFS Messages*.

The OM command timeout default of 300 seconds (5 minutes) might not be enough time for the online change phase to complete. You might need to specify a timeout value on the command based on the needs of the installation.

### **INITIATE OLC return and reason codes**

A return and reason code is returned to OM by the INITIATE OLC command. The return and reason codes for all IMS commands that are routed from OM are defined in the DFSCMDRR macro in the IMS.SDFS MAC data set.

The INITIATE OLC command master usually performs the online change phase locally first. If the online change phase fails locally, the command master usually skips sending the online change phase to the other IMS systems, sets a completion

code for each other IMS indicating that the online change phase was not attempted, and terminates command processing. However, if the INITIATE OLC PHASE(COMMIT) command fails on the local IMS because of work in progress, the command master still sends the commit phase 1 to the other IMS systems. The purpose is to report work in progress for all the IMS systems in the IMSplex, to facilitate completion of work in progress.

In a mixed IMSplex, you might have some IMS systems that support a particular version of online change and some that do not. For the command to be considered successful, at least one IMS in the IMSplex must successfully perform the online change phase. If no IMS in the IMSplex supports the same version of online change, the command reason code indicates that none of the IMS systems performed the online change phase. If you enter a version of the INITIATE OLC PHASE(PREPARE) command that does not apply to any IMS in the IMSplex, you must terminate the online change by using a TERMINATE OLC command.

**Related reference:**

 Return and reason codes for commands to OM (Messages and Codes)

### **Handling errors for the INITIATE OLC command**

The INITIATE OLC command can result in errors that leave one or more of the IMS systems in the IMSplex in various online change states. Correct the error. Issue the QUERY MEMBER command and the QUERY OLC command to help you decide whether to terminate the online change by issuing TERMINATE OLC or to try the INITIATE OLC command again.

Before attempting a global online change, issue the QUERY OLC LIBRARY(OLCSTAT) SHOW(MODID) command to get the current modify ID. If the INITIATE OLC command fails, issue the QUERY OLC LIBRARY(OLCSTAT) SHOW(MODID) command again, to see if the modify ID is the same. If the modify ID increased by 1, the online change is considered to be successfully completed.

If the INITIATE OLC command fails, issue the QUERY MEMBER TYPE(IMS) SHOW(STATUS) command to display the online change state of all the IMS systems in the IMSplex. Evaluate the QUERY MEMBER TYPE(IMS) SHOW(STATUS) output to help you to determine what to do. For information about the QUERY MEMBER TYPE(IMS) SHOW(STATUS) output and the action to take, see “Displaying status or attribute information about IMS systems in the IMSplex” on page 456.

## **Changing or adding IMS.ACBLIB members online**

In an IMSplex environment, you can use the ACB member online change (OLC) function to add or change individual members of the ACB library, or the entire ACB library, and bring these new or changed members online without quiescing the IMSplex or refreshing the active ACB library.

Members that are not affected by the member OLC are not quiesced.

You can make online changes to the ACB library in IMSplex environments in which:

- A single IMS system comprises the IMSplex. In this configuration, IMS must use the OLCSTAT data set with a CSL that consists of an OM and SCI. An RM is not required.
- Multiple IMS systems comprise the IMSplex. In this configuration, each IMS system must use the same shared OLCSTAT data set, or each system must use

its own local OLCSTAT data set. The CSL must consist of an OM, an SCI, and an RM. Although a resource structure is recommended, it is not required.

Global online change for ACBLIB members can only be performed in an IMSplex.

You can specify whether the ACBLIB member is shared or dedicated in the IMSplex. The following example highlights how the ACBLIB member online change process can be coordinated among all IMS systems that share the OLCSTAT data set:

- OM chooses an IMS system to be the command master IMS.
- The command master IMS coordinates the member OLC process with other IMS systems sharing the OLCSTAT data set using RM.
- ACBSHR= must be specified as Y or N in the DFSCGxxx or DFSDFxxx PROCLIB member.
  - Y– Indicates that all of the IMS systems in the OLCSTAT are using the same active and inactive ACBLIB.
  - N– Indicates that each IMS in the OLCSTAT is using its own dedicated active and inactive ACBLIB.
- All IMS systems in the IMSplex that share the OLCSTAT data set must specify the value of ACBSHR=.
- The member OLC updates all the active ACBLIBs, whether or not they are shared.

The ACB member online change uses IMS type-2 commands only. The commands that are involved in performing an ACB library member online change are:

- INITITATE OLC PHASE(PREPARE)
- INITIATE OLC PHASE(COMMIT)
- TERMINATE OLC
- QUERY MEMBER TYPE(IMS)
- QUERY OLC SHOW(RSCLIST)

If you want to fall back to the previous version of the changed resource, perform a full online change process with a full library switch.

To perform an ACB member online change:

1. Define the IMS system with online change capability by removing local MODSTAT definitions and defining the OLCSTAT data set in the DFSCGxxx PROCLIB member.
2. Run the PSB generation to generate or update PSBLIB members.
3. Run the DBD generation to generate or update DBDLIB members.
4. Run the ACB generation to create an executable ACBLIB member in a staging ACB library.

The staging ACB library can be a clone or a subset of the active ACB library.

5. Optionally, create a backup copy of the active ACBLIB member for fallback purposes; you can use the Online Change Copy utility (DFSUOCU0) to do this.
6. Issue the INITIATE OLC command with the TYPE(ACBMBR) parameter to copy the updated or new member from the staging library to the active library.

If your ACBs are in 64-bit storage, the ACB members that are affected by the online change process are removed from the non-resident pool and deleted from the 64-bit pool. The next time an application is scheduled that needs the



new or updated ACB members, the members are read from the ACB library and copied to the non-resident pool for execution and are also copied to the 64-bit pool.

If you want to fall back to the previous version of the changed resource, perform a full online change process with a full library switch.

**Related reference:**

 INITIATE OLC command (Commands)

## Displaying status or attribute information about IMS systems in the IMSplex

The QUERY MEMBER command displays status or attribute information about IMS systems in the IMSplex, including online change status.

You can specify QUERY MEMBER only through the OM API. For syntax, environment, and usage information, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

### Displaying information about the OLCSTAT data set

The QUERY OLC LIBRARY command provides information about the OLCSTAT data set contents. For syntax, environment, and usage information, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

## Terminating an online change

The TERMINATE OLC command supports aborting a global online change across an IMSplex. The TERMINATE OLC command can also be shortened to TERM OLC.

A TERMINATE OLC command that aborts a global online change is similar to the /MODIFY ABORT command, except that a global command applies to all IMS systems in an IMSplex that are participating in the global online change.

OM sends the TERMINATE OLC command to an IMS in the IMSplex.

You can use TERMINATE OLC to abort an IMSplex-wide global online change initiated by a INITIATE OLC PHASE(PREPARE), before the online change is successfully committed with a INITIATE OLC PHASE(COMMIT).

Use TERMINATE OLC to abort an online change after an INITIATE OLC PHASE(COMMIT) failure that occurs before the OLCSTAT data set is updated. After the commit process has updated the OLCSTAT data set, the online change is considered to be successful and cannot be aborted.

TERMINATE OLC is not supported if local online change is enabled. The TERMINATE OLC command is rejected if the IMS to which the command is routed does not support global online change. If this occurs and there is an IMS that supports global online change, the user must route the command to a specific IMS that supports global online change.

You can specify TERMINATE OLC only through the OM API.

The TERMINATE OLC command can be processed by DB/DC, DBCTL, and DCCCTL environments.

The TERMINATE OLC command is invalid on the XRF alternate, RSR tracker, and FDR system.

The OM command timeout default of 300 seconds (5 minutes) might not be enough time for the online change phase to complete. You might need to specify a timeout value on the command based on the needs of the installation.

The command syntax for this command is defined in XML and is available to automation programs which communicate with OM. For more information about command syntax, environments, and usage, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

## Handling errors for global online change commands

When a global online change command results in an error, the action you perform depends on the state of the IMS systems in the IMSplex after the online change error has occurred.

The following table describes actions that you can perform when an online change command results in an error.

*Table 35. Handling errors for online change commands*

| IMSplex state after online change command error                                     | User action to perform                                                                              |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| One or more IMS systems in prepare complete state after a prepare failure.          | Issue the TERMINATE OLC command. Correct the problem before reissuing the command.                  |
| All IMS systems in prepare complete state.                                          | Issue the INIT OLC PHASE(COMMIT) or TERMINATE OLC command.                                          |
| A mix of IMS systems in prepare complete and commit phase 1 complete state.         | Correct the problem and try the commit again or issue the TERMINATE OLC command.                    |
| All IMS systems in commit phase 1 complete state before OLCSTAT data set updated.   | Correct the problem and try the commit again or issue the TERMINATE OLC command.                    |
| All IMS systems in commit phase 1 complete state after OLCSTAT data set updated.    | Correct the problem and try the commit again. The online change is committed and cannot be aborted. |
| Mix of IMS systems in commit phase 1 and commit phase 2 complete state.             | Correct the problem and try the commit again.                                                       |
| All IMS systems in commit phase 2 complete state.                                   | Correct the problem and try the commit again.                                                       |
| Mix of IMS systems in commit phase 2 complete state and not in online change state. | Correct the problem and try the commit again.                                                       |

## Global online change scenarios

The scenarios in this topic describe events that can occur during the processing of online change commands and explains how the IMS systems respond to the commands based on their current state.

### Scenario 1: a successful global online change where all IMS systems use MODBLKS=OLC

This scenario shows the output returned to the TSO SPOC as the result of the successful global online change command sequence.



1. The QUERY OLC LIBRARY(OLCSTAT) command is issued to show the active online change libraries, ACBLIBA, FMTLIBA, and MODBLKSA. The command also shows that the current modify ID is 1 and the OLCSTAT data set name is IMSTESTL.OMS.OLCSTAT. The IMS systems that are current with the online change libraries include SYS3, IMS2, and IMS3. These IMS systems will be permitted to perform a warm start or emergency restart.
2. The INIT OLC PHASE(PREPARE) TYPE(MODBLKS) OPTION(FRCABND,FRCNRML) is issued to prepare IMS systems for a MODBLKS online change. IMS3 is the prepare master because it builds all of the output lines. Because the FRCABND and FRCNRML options are specified when the online change command is issued, any IMS systems that are down are permitted to perform the online change. IMS2, SYS3, and IMS3 are in the OLCSTAT data set, so they will participate in the online change.
3. The INIT OLC PHASE(COMMIT) is issued to commit the online change. IMS2, IMS3, and SYS3 all successfully complete the online change. The commit master, IMS3, reports the new online change libraries, MODBLKSB, ACBLIBA, and FMTLIBA, and the new modify ID of 2.
4. Another QUERY OLC LIBRARY(OLCSTAT) SHOW(ALL) command is issued to show the new online change libraries and the new modify ID of 2.

The following sample output represents four responses to the four commands described in this scenario:

Response for: QUERY OLC LIBRARY(OLCSTAT) SHOW(ALL)

| MbrName | CC | Library | ACBLIB | FMTLIB | MODBLKS | Modid | DSName               | LastOLC | MbrList        |
|---------|----|---------|--------|--------|---------|-------|----------------------|---------|----------------|
| IMS3    | 0  | OLCSTAT | A      | A      | A       | 01    | IMSTESTL.OMS.OLCSTAT |         | SYS3,IMS2,IMS3 |

Response for: INIT OLC PHASE(PREPARE) TYPE(MODBLKS) OPTION(FRCABND,FRCNRML)

| MbrName | Member | CC | ACBLIB | FMTLIB | MODBLKS | Modid |
|---------|--------|----|--------|--------|---------|-------|
| IMS3    | IMS2   | 0  |        |        |         |       |
| IMS3    | SYS3   | 0  |        |        |         |       |
| IMS3    | IMS3   | 0  | A      | A      | A       | 1     |

Response for: INIT OLC PHASE(COMMIT)

| MbrName | Member | CC | ACBLIB | FMTLIB | MODBLKS | Modid |
|---------|--------|----|--------|--------|---------|-------|
| IMS3    | IMS2   | 0  |        |        |         |       |
| IMS3    | SYS3   | 0  |        |        |         |       |
| IMS3    | IMS3   | 0  | A      | A      | B       | 2     |

Response for: QUERY OLC LIBRARY(OLCSTAT) SHOW(ALL)

| MbrName | CC | Library | ACBLIB | FMTLIB | MODBLKS | Modid | DSName               | LastOLC | MbrList                |
|---------|----|---------|--------|--------|---------|-------|----------------------|---------|------------------------|
| IMS3    | 0  | OLCSTAT | A      | A      | B       | 02    | IMSTESTL.OMS.OLCSTAT |         | MODBLKS SYS3,IMS2,IMS3 |

### Scenario 1a: a successful global online change where one IMS uses MODBLKS=DYN

The IMS that uses DRD does not participate in the global online change. You must use DRD on that IMS to change the resource that you are changing in the other IMS systems that participate in the global online change.

This scenario involves a global online change for an IMSplex that has four IMS systems with the following attributes:

- Three IMS systems that use the online change process for the resources in the MODBLKS data set. They have MODBLKS=OLC specified in the DYNAMIC\_RESOURCES section of the DFSDFxxx IMS.PROCLIB member.
- One IMS that uses dynamic resource definition (DRD). It has MODBLKS=DYN specified in the DYNAMIC\_RESOURCES section of the DFSDFxxx IMS.PROCLIB member.

The following list illustrates a successful global online change command sequence:

1. The QUERY OLC LIBRARY(OLCSTAT) command is issued to show the active online change libraries, ACBLIBA, FMTLIBA, and MODBLKSA. The command also shows that the current modify ID is 1 and the OLCSTAT data set name is IMSTESTL.OMS.OLCSTAT. The IMS systems that are current with the online change libraries include SYS3, IMS2, and IMS3. These IMS systems will be permitted to perform a warm start or emergency restart.
2. The INIT OLC PHASE(PREPARE) TYPE(MODBLKS) OPTION(FRCABND,FRCNRML) is issued to prepare IMS systems for a MODBLKS online change. IMS3 is the prepare master because it builds all of the output lines. Because the FRCABND and FRCNRML options are specified when the online change command is issued, any IMS systems that are down are permitted to perform the online change. IMS2, SYS3, and IMS3 are in the OLCSTAT data set, so they will participate in the online change.
3. The INIT OLC PHASE(COMMIT) is issued to commit the online change. IMS2, IMS3, and SYS3 all successfully complete the online change. The commit master, IMS3, reports the new online change libraries, MODBLKSB, ACBLIBA, and FMTLIBA, and the new modify ID of 2.
4. Another QRY OLC LIBRARY(OLCSTAT) SHOW(ALL) command is issued to show the new online change libraries and the new modify ID of 2.

The following sample output represents 4 responses to the 4 commands described in this scenario:

Response for: QRY OLC LIBRARY(OLCSTAT) SHOW(ALL)

| MbrName | CC | Library | ACBLIB | FMTLIB | MODBLKS | Modid | DSName               | LastOLC | MbrList        |
|---------|----|---------|--------|--------|---------|-------|----------------------|---------|----------------|
| IMS3    | 0  | OLCSTAT | A      | A      | A       | 01    | IMSTESTL.OMS.OLCSTAT |         | SYS3,IMS2,IMS3 |

Response for: INIT OLC PHASE(PREPARE) TYPE(MODBLKS) OPTION(FRCABND,FRCNRML)

| MbrName | Member | CC | ACBLIB | FMTLIB | MODBLKS | Modid |
|---------|--------|----|--------|--------|---------|-------|
| IMS3    | IMS2   | 0  |        |        |         |       |
| IMS3    | SYS3   | 0  |        |        |         |       |
| IMS3    | IMS3   | 0  | A      | A      | A       | 1     |

Response for: INIT OLC PHASE(COMMIT)

| MbrName | Member | CC | ACBLIB | FMTLIB | MODBLKS | Modid |
|---------|--------|----|--------|--------|---------|-------|
| IMS3    | IMS2   | 0  |        |        |         |       |
| IMS3    | SYS3   | 0  |        |        |         |       |
| IMS3    | IMS3   | 0  | A      | A      | B       | 2     |

Response for: QRY OLC LIBRARY(OLCSTAT) SHOW(ALL)

| MbrName | CC | Library | ACBLIB | FMTLIB | MODBLKS | Modid | DSName               | LastOLC | MbrList                |
|---------|----|---------|--------|--------|---------|-------|----------------------|---------|------------------------|
| IMS3    | 0  | OLCSTAT | A      | A      | B       | 02    | IMSTESTL.OMS.OLCSTAT |         | MODBLKS SYS3,IMS2,IMS3 |

## Scenario 2: online change prepare command is rejected because an IMS is down

In this online change scenario, an INITIATE OLC PHASE(PREPARE) TYPE(ALL) command was issued and routed to IMS A. The INITIATE OLC PHASE(PREPARE) command is rejected because two IMS systems are down (IMS B abends and IMS C shuts down normally), but no FRCABND or FRCNRML options were specified when the command was issued.

The following table identifies the command sequence and codes returned during this scenario:

Table 36. Scenario 2: command sequence and codes for failed online change prepare command

| Action or event                   | IMS A completion code | IMS B completion code | IMS C completion code | Command return code | Command reason code |
|-----------------------------------|-----------------------|-----------------------|-----------------------|---------------------|---------------------|
| INIT OLC PHASE(PREPARE) TYPE(ALL) | 2                     | B2                    | B2                    | C                   | 3004                |

Table 36. Scenario 2: command sequence and codes for failed online change prepare command (continued)

| Action or event                                           | IMS A completion code | IMS B completion code | IMS C completion code | Command return code | Command reason code |
|-----------------------------------------------------------|-----------------------|-----------------------|-----------------------|---------------------|---------------------|
| INIT OLC PHASE(PREPARE) TYPE(ALL) OPTION(FRCABND,FRCNRML) | 0                     | B2                    | B2                    | 4                   | 100C                |
| INIT OLC PHASE(COMMIT)                                    | 0                     | B2                    | B2                    | 4                   | 100C                |
| INIT OLC PHASE(PREPARE) TYPE(ALL)                         | 0                     |                       |                       | 0                   | 0                   |
| INIT OLC PHASE(COMMIT)                                    | 0                     |                       |                       | 0                   | 0                   |

1. An INITIATE OLC PHASE(PREPARE) TYPE(ALL) command was issued. IMS A is the command master. IMS A detects that IMS B and IMS C are down and that no FRCABND or FRCNRML keywords were specified, so it rejects the command. IMS A returns a completion code of 2 (not done) for itself; IMS A also returns a completion code of B2 (IMS state error) for IMS B and IMS C. IMS A sets an overall return code of C and an overall reason code of 3004 (failed for all IMS systems).
2. Another INITIATE OLC PHASE(PREPARE) TYPE(ALL) command was issued, this time with FRCABND and FRCNRML. These options force the online change even though IMS systems are down. IMS A is again the command master. IMS A detects that IMS B and IMS C are down, but determines that this is acceptable. IMS A sets a completion code of B2 (IMS state error) for IMS B and IMS C. IMS B and IMS C do not perform the prepare phase, because they are down. IMS A performs the prepare phase locally and returns a completion code of 0 for itself. IMS A then sets an overall return code of 4 and a reason code of 100C (failed for some IMS systems for acceptable reasons).
3. An INIT OLC PHASE(COMMIT) command is issued; the commit succeeds. IMS A updates the OLCSTAT data set with the new information. IMS A includes a record for itself, but omits a record for IMS B and IMS C. IMS B and IMS C are no longer current with the online change libraries. IMS B and IMS C will have to cold start because the last online change included MODBLKS. If an IMS is down during a MODBLKS online change, it has to cold start.
4. Another INITIATE OLC PHASE(PREPARE) command is issued. This time, IMS B and IMS C do not participate because they are not current with the online change libraries. Only IMS A participates in the global online change.
5. An INIT OLC PHASE(COMMIT) command is issued. IMS B and IMS C do not participate because they are not current with the online change libraries. Only IMS A commits the global online change. IMS A sets a completion code of 0 for itself and an overall return code and reason code of 0.

### Scenario 3: commit phase 1 fails due to work in progress

In this online change scenario, commit phase 1 fails due to work in progress. However, the INITIATE OLC PHASE(COMMIT) command eventually succeeds on all of the IMS systems, which results in overall return and reason codes of 0.

The following table identifies the command sequence and codes returned during this scenario:

Table 37. Scenario 3: command sequence and codes for failed commit phase 1

| Action or event                    | IMS A completion code | IMS B completion code | IMS C completion code | Command return code | Command reason code |
|------------------------------------|-----------------------|-----------------------|-----------------------|---------------------|---------------------|
| INIT OLC PHASE(PREPARE) TYPE(ALL)  | 0                     | 0                     | 0                     | 0                   | 0                   |
| INIT OLC PHASE(COMMIT)             | B1                    | B1                    | 4                     | C                   | 3004                |
| INIT OLC PHASE(COMMIT)             | B1                    | 4                     | 4                     | C                   | 3004                |
| /DISPLAY MODIFY and other commands | none                  | none                  | none                  | none                | none                |
| INIT OLC PHASE(COMMIT)             | 0                     | 0                     | 0                     | 0                   | 0                   |

1. The INITIATE OLC PHASE(PREPARE) command is issued and succeeds on all of the IMS systems. IMS A is the command master. Each IMS performs the prepare phase locally and sets a completion code of 0. IMS A sets an overall return code and reason codes of 0.
2. The INITIATE OLC PHASE(COMMIT) command is issued, and IMS A is the command master. IMS A attempts commit phase 1 locally, but the commit phase fails due to work in progress. IMS A sets a completion code of B1 (resource state) for itself. IMS A continues processing and sends the commit phase 1 to the other IMS systems that need to report any work in progress. IMS C successfully completes commit phase 1, but IMS B has work in progress. IMS A sets a completion code of 4 (incomplete) for IMS C and quits the commit process at this point due to errors.
3. Another INITIATE OLC PHASE(COMMIT) command is issued; IMS A is the master again. IMS A again attempts commit phase 1 locally, which fails due to work in progress. IMS A sends commit phase 1 to IMS B and IMS C. IMS B successfully completes commit phase 1. IMS C returns a completion code indicating that it is already in the correct state. IMS A quits the commit due to errors and sets the completion code to 4 (incomplete) for IMS B and IMS C. IMS A sets an overall return code of C and a reason code of 3004 (failed for all IMS systems).
4. /DISPLAY MODIFY and other commands are issued to resolve the work in progress that might cause the commit to fail.
5. A final INITIATE OLC PHASE(COMMIT) command is issued. All of the IMS systems successfully complete the commit phases 1, 2, and 3 and set completion codes of 0 for themselves. The overall return code and reason code is 0. IMS A sets an overall return code and reason code of 0.

#### Scenario 4: online change prepare command fails due to a timeout error

In this scenario, the online change prepare command fails due to a timeout error. Because one IMS never receives the prepare phase, the command times out while waiting for a response. Online change must be aborted after a prepare failure.

The following table identifies the command sequence and codes returned during this scenario:

Table 38. Scenario 4: command sequence and codes for failed prepare command due to timeout error

| Action or event                       | IMS A completion code | IMS B completion code | IMS C completion code | Command return code | Command reason code |
|---------------------------------------|-----------------------|-----------------------|-----------------------|---------------------|---------------------|
| INIT OLC PHASE(PREPARE) TYPE(MODBLKS) | 0                     | 0                     | 91                    | C                   | 3000                |

Table 38. Scenario 4: command sequence and codes for failed prepare command due to timeout error (continued)

| Action or event                       | IMS A completion code | IMS B completion code | IMS C completion code | Command return code | Command reason code |
|---------------------------------------|-----------------------|-----------------------|-----------------------|---------------------|---------------------|
| INIT OLC PHASE(PREPARE) TYPE(MODBLKS) |                       |                       |                       | 10                  | 4110                |
| INIT OLC PHASE(COMMIT)                | 2                     | 2                     | B2                    | C                   | 3004                |
| TERMINATE OLC                         | 0                     | 0                     | 3                     | 4                   | 100C                |
| INIT OLC PHASE(PREPARE) TYPE(MODBLKS) | 0                     | 0                     | 0                     | 0                   | 0                   |
| INIT OLC PHASE(COMMIT)                | 0                     | 0                     | 0                     | 0                   | 0                   |

1. In this scenario, IMS A is the command master, and the INITIATE OLC PHASE(PREPARE) TYPE(MODBLKS) command is issued; IMS A is the command master. SCI goes down before IMS C receives the prepare phase. IMS A and IMS B successfully complete the prepare phase locally and set completion codes of 0 for themselves. The RM process step response times out, and RM returns the process step results to IMS A. IMS A sets a completion code of 91 (time out) for IMS C. IMS A rejects the prepare command because of the timeout error. After the failure, IMS A and IMS B remain in an online change prepare state, and IMS C is not in an online change state. IMS A sets an overall return code of C and reason code of 3000 (prepare worked for some of the IMS systems). Because IMS C does not do the prepare and a prepare failure occurred, the online change must be aborted.

In this situation, you should issue QUERY MEMBER TYPE(IMS) to see the state of the IMS systems to determine what to do.

2. Another INIT OLC PHASE(PREPARE) TYPE(MODBLKS) command is issued. However, because online change must be aborted after a prepare failure, this prepare command also fails. IMS A, the command master, detects that it is already in a prepare state and rejects the command with a return code of 10 and a reason code of 4110. (The INIT OLC command is rejected because the command does not apply to the online change state of the command master.)
3. An INIT OLC PHASE(COMMIT) command is issued. The installation mistakenly tries to commit the online change after the prepare failure. IMS A is once again the command master. IMS A performs commit phase 1 locally, then sends commit phase 1 to IMS B and IMS C. IMS A and IMS B successfully perform commit phase 1 locally. IMS C rejects the commit phase 1 because it is not in an online change prepare state. IMS C returns a completion code of B2 (IMS state error). IMS A detects the error and sets the completion code to 2 (incomplete) for IMS A and IMS B. IMS A also returns a completion code of B2 (IMS state error). IMS A then terminates commit processing. The command fails with a return code of C and a reason code of 3004 (no IMS systems successfully completed command).
4. A TERMINATE OLC command is issued to abort the global online change. IMS A and IMS B successfully complete the abort. Because IMS C is not in an online change state, it returns a completion code of 3 (IMS already in requested online change state). IMS A, the command master, treats this as acceptable, and returns an overall return code of 4 and a reason code of 100C (successful, but not applicable to some IMS systems).
5. An INIT OLC PHASE(PREPARE) TYPE(MODBLKS) command is issued. The prepare succeeds.
6. An INIT OLC PHASE(COMMIT) command is issued. The commit succeeds.

---

## Performing a cold start

Under certain circumstances—for example, if an IMS goes down during an online change—you might need to perform a cold start. The process you follow to perform a cold start varies depending on whether the IMS uses local or global online change.

For general information about cold start, see “Cold start” in *IMS Version 11 Operations and Automation*.

### Local online change and cold start

Using local online change affects all types of IMS system definitions. IMS online subsystems place control blocks in the IMS.MODBLKS staging library. Before you perform a cold start, copy the MODBLKS library to either the IMS.MODBLKSA or the IMS.MODBLKSB data set.

If you copy these data sets to the inactive library, update the IMS.MODSTAT data set to change the IMS.MODBLKS data set DD name. If you copy to the active library, you cannot, in case of problems, return to the last system environment before the change.

### Global online change and cold start

**Attention:** Typically, an IMS that was down during a global online change needs to be cold started when it comes back up. A cold start might be required to prevent restart from processing log records against the current online change libraries, potentially causing restart to fail or resulting in a severe error later.

Before the first IMS in the IMSplex cold starts the first time, you must initialize the OLCSTAT data set. To initialize the OLCSTAT, use the Global Online Change utility (DFSUOLC0).

**Important:** Use the Global Online Change utility with caution, so you do not inadvertently destroy valid OLCSTAT data set contents. If you do destroy valid OLCSTAT data set contents, global online change and initialization of additional IMS systems fail until the OLCSTAT data set is re-initialized.

You should establish an OLCSTAT data set recovery procedure to deal with the loss of the OLCSTAT data set. After every successful global online change, record the modify ID, the active online change library suffixes, and the list of IMS systems that are current with the online change libraries. If the OLCSTAT data set is destroyed, run the initialize function of the Global Online Change utility with the saved data to re-initialize the OLCSTAT data set.

**Related reading:** For more information about the Global Online Change utility, see *IMS Version 11 System Utilities*.

## Determining when to perform a cold start

Perform a cold start when an IMS was down during the last global online change and its restart type conflicts with the last online change type or when an IMS was down for two or more global online changes.

Perform a warm start when an IMS was down during the last global online change and the restart does not conflict with the last global online change.



For more information about the types of restart commands permitted for specific online change types, see “Performing a cold start” on page 463.

*Table 39. Restart commands permitted for specific online change types*

| Online change type | Restart commands permitted                                        |
|--------------------|-------------------------------------------------------------------|
| ACBLIB             | /NRE CHECKPOINT 0<br>/ERE COLDBASE                                |
| ALL                | /NRE CHECKPOINT 0                                                 |
| FORMAT             | /NRE CHECKPOINT<br>/NRE<br>/ERE<br>/ERE COLDCOMM<br>/ERE COLDBASE |
| MODBLKS            | /NRE CHECKPOINT 0                                                 |

## Cleaning up after a failed restart attempt

If an IMS is down during global online change and your restart command is rejected because the IMS requires a cold start, you might need to perform one or more cleanup steps to ensure a successful cold start.

To clean up after a failed restart attempt:

1. Issue a DBRC LIST.RECON command to check for an IMS subsystem record or any open OLDSs.
2. If a subsystem record exists, issue a DBRC DELETE.SUBSYS command to clean up the DBRC subsystem record for that IMS.
3. For any open OLDSs, issue a DBRC GENJCL.CLOSE command. This command enables DBRC to read the RECON data set and determine which OLDS to close for that IMS.

---

## Chapter 27. Printing output with IMS Spool API

This topic provides design and operational advice for the use of the IMS DL/I Spool application programming interface (API) and provides details for using the IMS Spool API to increase the access of IMS application programs to advanced printing capabilities.

IMS provides an expansion of the DL/I application programming interface that allows application programs to interface directly to JES and to create print data sets on the JES spool. These print data sets are then made available to print managers and spool servers to serve the needs of the application program.

---

### Design and operational considerations

The following topic introduces IMS support of IMS Spool API.

#### Native IMS terminal support

IMS supports a set of terminals and printers that are associated with the IMS control region through the IMS systems generation process and ETO.

For most environments, a message destined for IMS-supported terminals is placed in the IMS message queue for intermediate storage. When the terminal can receive the message, IMS retrieves the message from the IMS message queue and sends the message to the device. If the VTAM session is lost during message transmission, or if a transmission error occurs, IMS places the message back on the IMS message queue for transmission at a later time. Because the IMS message queue is a recoverable resource, IMS can provide message integrity and recover from media failures associated with the IMS message queue.

The following figure shows a simplified diagram of the IMS environment. The figure shows terminal devices controlled by the IMS control region, the IMS message queue, and a message processing region. The DL/I pre-processing routines and the services of the IMS control region provide both terminal and database services. An application program sending output messages to a printer can identify the printer through the DL/I CHNG call and insert messages to the printer through the DL/I ISRT call. IMS places these messages in the IMS message queue until they are committed by the application program. After the calls are committed, IMS delivers the messages to the designated printer.






Figure 47. IMS environment without Spool API

## Application requirements

IMS application programs can send messages to IMS native printers and to printers that are capable of printing high-quality text, special fonts, bars, and graphs. Application programs that create data streams for these printers and send data to these printers use the job application subsystem IMS Spool API. The IMS Spool API passes the print data sets to IBM's Advanced Function Printing (AFP) services or to an OEM print server based on data set class, remote destination, size (segment sizes of up to 32 KB are common), or other characteristics.

Applications creating JES print data sets from IMS application regions require reasonable performance, a simple interface, and flexibility in program usage. The IMS Spool API interface extensions attempt to provide for these needs. The DL/I application interface supports the creation of JES print data sets. The techniques for creating these JES print data sets are similar to the techniques used for sending messages to native IMS printers.

IMS uses several z/OS services to support IMS Spool API. These z/OS services introduce performance and availability considerations that IMS and its application programs have not had to deal with previously, ranging from additional processor requirements for data set options parsing to lack of sync point support by the JES subsystem. The IMS Spool API environment, shown in the following figure, is designed to address these issues.

Because the print data sets can be very large and because using the message queue as intermediate storage adds processing overhead without providing additional message integrity, most IMS Spool API execution is performed in the IMS dependent region. However, this asynchronous message technique also presents problems in relating application program errors that result in message delivery

problems.




Figure 48. IMS environment with Spool API

**Note:** Do not confuse the support for creating JES print data sets on the IMS Spool API with the IMS terminal support for spool lines defined as UNITYPE=SPOOL in the IMS systems generation process. Terminal support for spool lines is completely unrelated to support for JES print data sets.

## The IMS Spool API as a data manager

The IMS Spool API controls the input, execution, and output of jobs on the z/OS system. Because the IMS Spool API is at the JOB or time sharing (TSO) user boundary, JES creates special considerations for IMS environments where each print data set is a logical message and is not part of or associated with a job boundary.

A data manager associated with IMS provides the DL/I application programming interface, which provides temporary buffering of data with the sync point process. The sync point process commits the changes for a unit of work if the unit of work terminates normally, or backs out the changes if the unit of work terminates abnormally. The IMS Spool API does not use the sync point process, but it can back out the effect of changes made by an abending unit of work.

The IMS Spool API does not recover data as does the IMS Transaction Manager or the IMS Database Manager. Because the IMS Spool API does not log changes to the spool data set, data on the spool can be lost in a restart if the spool device finds data errors.

---

## Print data set characteristics

You can create a JES print data set by allocating a print data set by the SYSOUT= specification on a JCL statement or dynamically by allocating a print data set using the MVS dynamic allocation support for SVC 99.

When a print data set is allocated, the characteristics of the print data set can be defined or allowed to default.

Some examples of these print data set characteristics are FORMS, COPIES, DESTINATION, and WRITER. You can specify some of these output characteristics with the DD card, or you can associate the print data set with a set of print data set characteristics defined by an OUTPUT JCL statement. JES provides a way to dynamically define these print data set descriptors using system services known as dynamic output or SVC 99 processing. When used with dynamic allocation, dynamic output processing can associate system output or SYSOUT data sets with a set of dynamically built OUTPUT descriptors.

You can build print data set descriptors for IMS Spool API in three ways, each of which is discussed in the following topics.

**Related reading:** For more information on z/OS services for dynamic allocation for print data sets, see *z/OS MVS Programming: Authorized Assembler Services Guide*.

## The Change (CHNG) call

In most application programs, you can provide the print data set descriptors using the change (CHNG) call interface. The IMS Spool API uses this existing technology.

The CHNG call supports two optional parameters for IMS Spool API:

- The options list parameter points to a character string containing a list of options. The IAFP keyword identifies the CHNG call as a request for IMS Spool API services.
- The feedback area parameter points to an area through which IMS can return error information to the application program when mistakes are found in the options list.

If your application creates multiple data sets with the same print options, or is wait-for-input (WFI), use the SETO call to reduce the parsing impact.

The CHNG call provides the data set characteristics in one of the following ways:

- The call can provide the data set options directly by using the PRTO option.

When the PRTO option is used with the CHNG call, IMS activates z/OS services (SJF) to verify the print options and calls z/OS services for dynamic output to create the output descriptors that are used when the print data set is allocated. This is the simplest way for the application program to provide print data set characteristics. Using the PRTO option involves the most overhead, because parsing must occur for each CHNG call with the PRTO option.

Use the PRTO option when you cannot significantly improve application performance by using pre-built text units across multiple CHNG calls.

- The call can reference a set of pre-built text units for dynamic descriptors using the TXTU option.

Application programs that can reuse text units can achieve better performance by using the TXTU option. If the application program can manage the text units necessary for dynamic output, use the TXTU option to avoid parsing for many of the print data sets.

The application program builds the text unit in the necessary format within the application area and passes these text units to IMS.

The application program can provide the print options to IMS with a SETO call and provide a work area for the construction of the text units. Using the TXTU option, IMS can perform parsing and text unit construction so that the work area passed contains the text units necessary for dynamic output after a successful SETO call. Do not relocate this work area because it contains address-sensitive information.

- The call can refer to an OUTPUT JCL statement in the dependent region's JCL by using the OUTN option.

This is a simple way of providing print data set information if it can be used by the application program. The output JCL statements are referenced by the OUTN option in the options list on the CHNG call. When the OUTN option is used, IMS references the output JCL statement at dynamic allocation so that the IMS Spool API obtains the print data set characteristics from the output JCL statement.

## The Set Options (SETO) call

You can construct a print data set descriptor with the set options (SETO) call. The SETO call is especially useful to application programs with wait-for-input (WFI) execution. The SETO call reduces the overhead necessary to perform parsing and text unit construction of the output descriptors for a data set.

If the application program needs to reuse a set of descriptors during the scheduling of the program, the application program provides the print data set characteristics to the IMS Spool API through the SETO call. The SETO call parses the output options and builds the dynamic output text units in the work area provided by the application program. Once the application program is supplied the pre-built text units, these text units are used with a CHNG call through the TXTU parameter. They provide the print characteristics for the data set without incurring the overhead of parsing and text unit build.

It is not necessary to use the SETO option to pre-build the text units if they can be pre-built using some other programming technique.

## The Output DD statement

You can provide print data set descriptors by adding output JCL statements to the dependent region JCL. The application program passes the name of the output JCL

statement in the dependent region JCL using a CHNG call with the OUTN parameter. IMS can use the output JCL statement when the print data set is dynamically allocated.

If the output statement does not exist in the dependent region's JCL, dynamic allocation fails when the first insert is done. The application program receives a status code AX indicating the insert (ISRT) failed.

---

## Writing data to the IMS Spool API

You can write data to the IMS Spool API using the insert (ISRT) call or the purge (PURG) call.

You can backout print data sets using the ROLL call or the ROLB call.

Do not use the SETS call or the ROLS call to backout print data sets. These calls are not supported by the IMS Spool API.

### The Insert call

The standard insert (ISRT) call writes the data to the IMS Spool API using BSAM. Because the length of the data written to the spool can be up to 32760 bytes, the BSAM write is performed directly from the application program's buffer area.

If IMS finds that the user's I/O area is above the 16 MB line, IMS moves the user's application program data to a work area below the line before performing the BSAM write.

If the application area is already below the line, the write can be done directly from the I/O area. If possible, the I/O area should be below the 16 MB line. However, you do not need to take unusual steps to place the I/O area below the line unless performance shows that such action is necessary. By writing the application program's buffer without first moving it, IMS prevents problems associated with moving large blocks of data. When BSAM is used, change the format of the I/O area so that it contains the BSAM block descriptor word (BDW).

**Related reading:** For more information on the format of the I/O area, see *IMS Version 11 Application Programming*.

### The PURG call

Use the purge (PURG) call with an express alternate PCB to release a print data set to the IMS Spool API for immediate printing. When the PURG call is issued, the print data set is closed and deallocated. If the PURG call is issued against an alternate PCB that is not marked as EXPRESS=YES, no action is taken for the print data set.

You can issue the PURG call either with or without an I/O area:

- When the PURG call is issued with an I/O area, it is treated as though it were two calls. The first function handles the PURG request, and the second function inserts the data provided by the I/O area.
- When the PURG call is issued against an alternate PCB that is not marked as EXPRESS=YES, the purge function is ignored. If an I/O area is included, the data is placed in the print data set.

The PURG call differs slightly with IMS Spool API when it is used with alternate PCBs that are generated as EXPRESS=YES:

- The current print data set is closed, deallocated, and sent to JES for printing.
- Status code of A3 is returned to the application program, indicating that the alternate PCB no longer contains a valid destination.

For information on the results of the PURG call, see the following table.

*Table 40. Results of PURG call*

| Type of PURG call | Express PCB                                                                                           | Non-express PCB                                |
|-------------------|-------------------------------------------------------------------------------------------------------|------------------------------------------------|
| With I/O area     | The print data set is released for printing by JES. PURG call receives status code of <b>A3</b> .     | The PURG call functions the same as ISRT call. |
| Without I/O area  | The print data set is released for printing by JES. PURG call receives status code of <b>blanks</b> . | No action is taken.                            |

## ROLL and ROLB calls

The ROLL and ROLB calls can be issued by an application program using the IMS Spool API functions. These calls result in backout (data set deletion) of any print data sets that have not been released to the IMS Spool API because of a PURG call to an express alternate PCB.

## SETS, SETU, and ROLS calls

The IMS Spool API does not support SETS, SETU, or ROLS calls. If you issue them, no action is taken by the IMS Spool API, and no special status codes are returned to show that the SETS, SETU, or ROLS call was issued by an application program.

No restrictions exist on the use of SETS, SETU, and ROLS calls, because these calls can be used by the application program outside the processing of print data sets.

---

## Special considerations—descriptors allowed

The output descriptors allowed by the IMS Spool API are those supported by the TSO OUTDES command of the z/OS system under which IMS is executing at the time that parameter validation is performed.

**Related reading:** For more information on output descriptors, see *z/OS TSO/E Command Reference*.

The initialization parameters you use with the IMS Spool API can have an impact on the parameters that you use to define print data set characteristics.

## Controlling print data sets

Application programs have some control over the IMS Spool API files IMS creates for them. This control applies to the data streams generated by the application program as well, because IMS is not sensitive to this data.

You can embed control information within the data stream itself, or you can specify the IMS Spool API file options in an application program DL/I CHNG call. These processing options are specified on a data set basis and are specified on the TSO OUTDES statement.

The IMS Spool API routes print data based on output scheduling. Generally, scheduling is determined by output class (CLASS parameter), printer destination (DEST parameter), FORM, and PRMODE.

IMS is a transparent delivery service. IMS dynamically creates IMS Spool API files based on data set processing options, writes the messages to them, and then deallocates them so that they can be processed by an IMS Spool API data set server (such as PSF).

To ensure proper disposition of print data sets, the z/OS operator might need to become involved under certain conditions, such as:

- IMS emergency restart
- Dependent region abnormal termination
- Dynamic deallocation failure for a print data set

Under such conditions, the z/OS operator must deal appropriately with the IMS Spool API data sets that represent in-doubt messages. IMS provides informational messages to help the z/OS operator if the application program requests the proper disposition of these data sets.

Review operational procedures for unprintable data sets at your installation. Many installations have jobs that delete print data sets after a given number of days. Make sure that these jobs do not delete IMS print data sets that should be printed.

## Express alternate PCBs

The IMS Spool API supports the use of express alternate PCBs. If the application program issues a purge (PURG) call against an alternate PCB that has been generated with EXPRESS=YES, the print data set is closed and deallocated. It is available for JES processing before the termination of the IMS unit of work.

**Related reading:** For more information on using the PURG call with an express alternate PCB, see “The PURG call” on page 470.

## XRF environments

The IMS Spool API operates within XRF systems with no special support. During an XRF takeover, the IMS Spool API is affected as if a failure and normal restart occurred. If partial print data sets exist when the takeover occurs, and IMS abnormal termination processing does not execute on the primary system, the partial print data sets can be made available to IMS Spool API for printing.

**Related reading:** For a description of the proper procedure for these partial print data sets, see *IMS Version 11 Application Programming*.

---

## Understanding allocation errors

The IMS Spool API interface dynamically allocates the print data set only after data is actually inserted to the data set. This reduces overhead and simplifies cleanup for abending transactions.

Occasionally, errors can occur during dynamic allocation that are the results of incorrect data set print options supplied with the CHNG or SETO call. The data set print options can be parsed during the processing of the CHNG and SETO calls. Destinations can only be validated during dynamic allocation.

If any of the print options are incorrect and dynamic allocation fails when the first insert is done for the data set, the ISRT call receives a status code of AX. The IMS Spool API code provides the following to the application program:

- A status code
- An error message DFS0013E
- A diagnostic log record (67D0)

The error message shows the type of service that is activated and the return and reason codes that are responsible for the error. For example, a common failure is indicated by reason code 046C: Remote work station not defined to job entry subsystem. You see this reason code if you select an invalid destination or you select integrity option 2 (non-selectable destination) when the destination of IMSTEMP has not been defined to JES. Specifying an invalid destination in the destination name parameter of the call results in a dynamic deallocation error when IMS deallocates the print data set.

Some of the services indicated by the error message include:

**DYN** z/OS dynamic allocation (SVC99)  
**OPN** z/OS data set open  
**OUT** z/OS dynamic output descriptors build (SVC109)  
**UNA** z/OS dynamic deallocation (SVC99)  
**WRT** z/OS BSAM write

**Related reading:** If the service is for dynamic allocation or deallocation, or for dynamic output descriptor build, see *z/OS MVS Programming: Authorized Assembler Services Guide*.





---

## Chapter 28. Accessing external systems from within IMS

IMS application programs can access external systems to use data that resides in those external system or to interact with business logic that resides outside of the IMS environment.

IMS application programs that want to use external subsystem data can use either or both of the following facilities:

- The External Subsystem Attach Facility (ESAF), which can be used by MPP, BMP, and IFP application programs, can access external subsystem data, such as DB2 for z/OS databases
- The DB2 Recoverable Resource Manager Services Attach Facility (DB2RRMS), which can be used by JMP and JBP applications that access DB2 for z/OS databases

You can specify both ESAF and DB2RRMS access in one SSM PROCLIB member.

IMS application programs that want to interact with business logic that resides outside of the IMS environment can do so using the IMS callout function. The callout requests can be either synchronous or asynchronous. The target of the callout request can be:

- Any generic web service, after going through IMS Enterprise Suite SOAP Gateway.
- A message-driven bean, an Enterprise JavaBeans component, a Java EE application, or web service, after going through IMS TM Resource Adapter.
- A user-written IMS Connect TCP/IP application or a vendor-supplied solution that uses TCP/IP and the IMS Connect protocol to retrieve callout requests.

### **Related reading:**

- For more information about installing and defining external subsystems to IMS, see:
  - “Accessing multiple external subsystems” in *IMS Version 11 Communications and Connections*.
  - “Specifying the SSM= EXEC parameter” in *IMS Version 11 System Definition*.
- For more information about the software requirements of the callout function, see *IMS Version 11 Release Planning*. For more information about how application programs can use the callout function, see *IMS Version 11 Application Programming*.



---

## Part 3. DBRC administration

This topic provides an overview of the IMS Database Recovery Control facility (DBRC) and a discussion of the interactions between IMS and DBRC.



---

## Chapter 29. Overview of DBRC

A feature of the IMS Database Manager that facilitates easier recovery of IMS databases, DBRC maintains information that is required for database recoveries, generates recovery control statements, verifies recovery input, maintains a separate change log for database data sets, and supports sharing of IMS databases and areas by multiple IMS subsystems.

DBRC is an integral part of IMS. IMS relies on DBRC to record and manage information about many items. DBRC keeps this information in a set of VSAM data sets that are collectively called the recovery control (RECON) data set and advises IMS (based on the information in the RECON data set) about how to proceed.

Specifically, DBRC:

- Helps you ensure IMS system and database integrity by recording and managing information associated with the logging process.
- Assists IMS in the restart process by notifying IMS which logs to use for restart.
- Assists IMS to allow or prevent access to databases in data-sharing environments by recording and managing database authorization information.
- Facilitates database and log recovery by:
  - Controlling the use and integrity of the information in the logs.
  - Recording and maintaining information about the databases and logs in the RECON data set.
  - Generating and verifying the Job Control Language (JCL) for various IMS utility programs.
- Supports Extended Recovery Facility (XRF) by identifying (in the RECON data set) if the subsystem is XRF capable.
- Supports Remote Site Recovery (RSR) by containing the RSR complex definition in the RECON data set and providing other services associated with controlling RSR.
- Supports IMSplexes by:
  - Notifying all DBRCs in the same IMSplex within the same DBRC sharing group when one of the DBRCs performs a RECON data set reconfiguration.
  - Allowing parallel access to the RECON data set.

DBRC also plays a key role in managing the log data needed to restart and recover IMS online subsystems.

Most IMS configurations require DBRC, including:

- Online configurations: DB/DC, DCCTL, or DBCTL
- Data-sharing environments, including IMSplexes
- Configurations that use Extended Recovery Facility (XRF)
- Remote Site Recovery (RSR)

**Attention:** DBRC is not required for IMS batch jobs and for some offline utilities. However, if batch jobs and utilities that access registered databases are allowed to run without DBRC, the recoverability and integrity of the databases could be lost. Even if your configuration does not require the use of DBRC (such as in a non-data sharing, non-RSR batch environment), you can simplify your recovery process by using DBRC to supervise recovery and protect your databases.

**Related reading:**

- *IMS Version 11 Operations and Automation* provides detailed descriptions of recovery procedures with and without DBRC.
- Chapter 32, “How DBRC helps in recovery,” on page 493 further describes how DBRC helps with database recovery.

**Related concepts:**

“Database Recovery Control facility” on page 8

“Changing OLDS characteristics” on page 109

“Database Recovery Control and recovery” on page 113

---

## DBRC components

DBRC includes the RECON data sets, the Database Recovery Control utility (DSPURX00), and skeletal JCL.

### RECON data sets

DBRC stores recovery-related information in the RECON data sets.

**Recommendation:** Define three VSAM KSDSs for the RECON data sets when you install DBRC. The first two data sets are active data sets; the third one is a spare. The second active data set is a copy of the first. For most purposes, you can think of the two active RECON data sets as if they were a single data set, the RECON data set.

**Related reading:** These data sets are described in detail in Chapter 41, “Initializing and maintaining the RECON data sets,” on page 535.

### Database Recovery Control utility (DSPURX00)

Use the Database Recovery Control utility to issue DBRC commands.

The DBRC commands allow you to perform all of the following tasks:

- List the information in the RECON data set
- Update the information in the RECON data set
- Use the information in the RECON data set to generate jobs for the IMS utilities

### Skeletal JCL

DBRC uses partitioned data set (PDS) members as input models (or templates) for generating input for some of the recovery utilities. These PDS members are distributed with IMS and are called skeletal JCL.

DBRC uses the skeletal JCL, information from the RECON data set, and instructions from a GENJCL command to generate the JCL and control statements that are needed to correctly run some of the recovery utilities. Modify the skeletal JCL to reflect your installation's system configuration.

**Related reading:** For more information about skeletal JCL, see *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*.

IMS provides PDS members that contain skeletal JCL statements. These PDS members are called skeletal JCL execution members and can be found in the IMS.SDFSISRC target library.

The IMS Installation Verification Program (IVP), customizes the skeletal JCL execution members and places the customized members into the IMS PROCLIB library. DBRC uses the IMS PROCLIB members to generate jobs (JCL and control statements) for the IMS utilities listed in *IMS Version 11 Database Utilities*. There is also a skeletal JCL execution member, JOBJCL, that produces a JOB statement.

**Related reference:**

 Database Recovery Control utility (DSPURX00) (System Utilities)

---

## DBRC tasks

There are several tasks that are automatically performed by DBRC and a few additional ones that you initiate with a command or request to DBRC.

DBRC automatically performs the following tasks:

- Records and manages information about logs for IMS
- Records recovery information in the RECON data set
- Verifies that database utilities have the correct input
- Controls the recovery of databases that are registered with DBRC
- Controls the access authorization information for the control and serialization of shared databases

You can perform the following tasks and more when you initiate them by passing commands or requests to DBRC:

- Start or stop the DBRC application programming interface (API)
- Record recovery information in the RECON data set
- Generate JCL for various IMS utilities and generate user-defined output
- List general information in the RECON data set
- Gather specific information from the RECON data set

**Related reading:**

- See Chapter 30, “Recording and controlling log information,” on page 485 for additional information about DBRC's logging support.
- See Chapter 34, “Supporting data sharing,” on page 499 for additional information about DBRC's data sharing support.
- See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for detailed information about the DBRC commands.
- See *IMS Version 11 System Programming APIs* for detailed information about the DBRC API requests.

---

## DBRC commands and API requests

Use DBRC commands or DBRC application programming interface (API) requests to obtain services from DBRC.



## DBRC commands

Use DBRC commands to generate a listing of the information in the RECON data set; add to, change, and delete information in the RECON data set, and generate the JCL and the control statements necessary to run the IMS utilities that are used in database recovery.

The following is a list of the DBRC batch commands, along with information about how they are used:

- Use the BACKUP.RECON command to create a backup copy of the RECON data set.
- Use the CLEANUP.RECON command to delete old or expired recovery-related information and log information from the RECON data set.
- Use the CHANGE commands to modify information in the RECON data set.
- Use the DELETE commands to delete information from the RECON data set.
- Use the GENJCL commands to generate jobs for the various IMS recovery utilities.
- Use the INIT commands to make the following changes to a RECON data set:
  - Create DB groups
  - Create DBDS groups
  - Create recovery groups
  - Register databases
  - Create image copy records that are available for use during subsequent runs of the supported image copy utilities
  - Create change accumulation data sets that are available for future use by the Change Accumulation utility
  - Create entries that defines an area data set (ADS) that belongs to a Fast Path area
  - Create CA groups
  - Register DBDSs or DEDB areas
  - Define global service groups
  - Register HALDB partitions
  - Initialize a RECON data set
- Use the LIST commands to produce a formatted printout of all or selected parts of the RECON data set.
- Use the NOTIFY commands to add to the RECON data set information that is normally written there automatically.
- Use the RESET.GSG command after an unplanned RSR takeover to remove obsolete recovery information about RSR-covered databases and areas from the original active site RECON data sets.

You can issue DBRC commands by embedding them in the JCL for running the Database Recovery Control utility (DSPURX00). This utility can process these commands while running either in a batch environment or as a TSO foreground program.

You can also issue the DBRC commands (with the exception of the LIST commands) as a parameter of the COMMAND API request (see *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*).

Certain DBRC commands can be issued from an online IMS environment (/RMxxxxxx commands). The online DBRC commands are:

- /RMCHANGE
- /RMDELETE
- /RMGENJCL
- /RMINIT
- /RMLIST
- /RMNOTIFY

## DBRC API requests

Use the DBRC API requests to obtain services from DBRC.

**Related reading:** For detailed information about the DBRC application programming interface and the DBRC API requests, see *IMS Version 11 System Programming APIs*.

**Related reference:**

 Database Recovery Control commands (Commands)



## Chapter 30. Recording and controlling log information

As IMS operates, it constantly records its activities in the IMS logs. After you have initialized DBRC, it participates in the IMS logging process by recording and controlling information about logging activities in the RECON data set.

The IMS logs contain information about:

- IMS startups and shutdowns
- Changes made to databases
- Transaction requests received and responses sent
- Application program initializations and terminations
- Application program checkpoints
- IMS system checkpoints

The IMS logs consist of the:

- Write-ahead data set (WADS)
- Online log data set (OLDS)
- System log data set (SLDS)
- Recovery log data set (RLDS)

The following figure shows the log data sets and restart data set (RDS) that IMS produces and the RECON data set that DBRC creates and maintains.




Figure 49. Logs produced for recovery and restart

### **Related reading:**

- “Logging” on page 92 for detailed discussions of the IMS logging process.

- “DBRC log-related commands” on page 488

After you have initialized DBRC, it participates in the IMS logging process by recording and controlling information about logging activities. This information is recorded in the RECON data set. If you want DBRC to control the recovery of your database data sets (DBDSs), you must register them with DBRC.

DBRC automatically records many items in the RECON data set, including:

- Information about log data sets
- Information about database data sets
- Information about events, such as:
  - Database updates
  - Database authorizations
  - Creation of database image copies
  - Reorganizations of databases
  - Recoveries of databases
  - Archiving of an OLDS and creation of the corresponding SLDS and RLDS
  - Execution of the Log Recovery utility
  - Subsystem sign-on
- Definitions and information about events for Remote Site Recovery

IMS uses this information for restarting itself and for database recovery jobs (if the databases are registered with DBRC). DBRC also tracks the archiving requirements of the OLDS and, if requested, generates and submits the JCL for archiving jobs.

DBRC also provides unit-of-recovery management for all attached subsystems. DBRC provides information about these units of recovery for batch backout, dynamic backout, partial backout, and restart.

For logs produced by batch systems, you are not required to use DBRC. The advantage of using DBRC for batch jobs is that DBRC records information about all the log data sets that are produced by batch jobs and prevents batch update jobs from executing if you specify a dummy or null log data set.

**Attention:** If registered databases are updated without DBRC control, DBRC cannot correctly control recovery for the databases and database integrity can be jeopardized.

---

## Changing log records in the RECON data sets

Use the `CHANGE.PRILOG` or `CHANGE.SECLOG` commands to modify the information about OLDSs, RLDSs, or SLDSs in the log record of the RECON data set.

Update the log record to indicate when errors have occurred on the data sets or that volume serial numbers have changed. You do not normally need to use these commands.

---

## Archiving an OLDS

Run the Log Archive utility (DFSUARCO) to archive an OLDS to an SLDS so that IMS can reuse the OLDS. How frequently you should archive depends on the load on the subsystem and the number of log records written to the OLDSs.

The Log Archive utility always produces an SLDS. The SLDS contains all log records that are required for both database recovery and for online IMS restart.

You can ask the Log Archive utility to produce an RLDS in addition to an SLDS. The RLDS contains only those log records that are required for database recovery.

If you request an RLDS, information about the output RLDS data sets is recorded in the PRILOG record in the RECON data set and information about the SLDS data sets is recorded in the PRISLD record. If you do not request an RLDS, the same information about the SLDS data sets is recorded in both the PRILOG and PRISLD records.

If there is a secondary OLDS, or if you request that dual logs be produced from a single OLDS, the information about the secondary-log output is recorded in corresponding SECLOG and SECSLD records.

**Important:** Log data sets that are output from IMS batch jobs are technically SLDSs, but the information about them is recorded in the PRILOG and SECLOG records.

Run the Log Archive utility by issuing the GENJCL.ARCHIVE command. DBRC then determines which OLDSs are full, and generates the appropriate JCL.

**Related reading:** See member DFSVSMxx in *IMS Version 11 System Definition* for more information on the ARCHDEF statement and automatic archiving.

**Recommendation:** Whether you use automatic archiving or invoke archiving yourself, make sure the archive jobs run as quickly as possible. The online subsystem only reuses an OLDS after it has been archived. If the archive job is not run and all the OLDS become full, the online subsystem waits. One way to ensure that archive jobs run quickly is to use an initiator that runs with a fairly high priority and is not used by many other users. This ensures that the archive jobs do not remain on the internal reader queue for too long.

If DBRC has marked an OLDS in the RECON data set as having errors, the GENJCL function does not submit it for archiving. If one of a pair of OLDSs has been destroyed or is unavailable, you can choose to mark it in the RECON data set as having errors.

The following references point to where you can find more information about archiving log records.

**Related reading:**

- See *IMS Version 11 Operations and Automation* for more information about automatic, manual and custom archiving of log records.
- See *IMS Version 11 System Utilities* for more information about specifying entry points and running the Log Archive utility.
- Refer to *IMS Version 11 Exit Routines* for more information about the Log Archive and the Logger exit routines.

---

## DBRC log-related commands

You can use commands to perform DBRC log-related functions in your operational procedures.

The following commands perform the DBRC log-related functions:

- CHANGE.PRILOG
- CHANGE.RECON
- CHANGE.SECLOG
- DELETE.LOG
- GENJCL.ARCHIVE
- GENJCL.CLOSE
- LIST.LOG
- NOTIFY.PRILOG
- NOTIFY.SECLOG

In addition to the LIST.LOG command, you can use a Log or OLDS query API request to retrieve log-related information from the RECON data set. You can also use the COMMAND API request to add to, change, and delete information in the RECON data set.

**Related reading:** See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for more information about DBRC commands.

---

## Chapter 31. Recovering databases

The recovery process for IMS databases can include these three basic steps, although the details of the process can vary with the type of database to be recovered.

1. Restore the database to the most current image copy.  
**Attention:** Skip this step if an HALDB Online Reorganization is being used because all the database changes are on log data sets.
2. Use the log data sets (or change accumulation data sets) to restore changes made to the database since the image copy was made.

**Definition:** Change accumulation is the process of creating a compacted version of one or more IMS log data sets by eliminating records not related to recovery, and by merging multiple changes to a single segment into a single change. For more information about change accumulation, see “Log record change accumulation” on page 525.

3. Back out any incomplete changes.

When the recovery completes successfully, DBRC records information about the recovery in the RECON data set. If you performed a time-stamp recovery, DBRC records information about the range of omitted changes.

**Requirement:** If a time-stamp recovery is performed within a database allocation interval, you must immediately create an image copy to ensure a valid starting point for subsequent recovery attempts. DBRC then prevents the changes from being reapplied on any subsequent recovery.

The following figure illustrates a simple database recovery.




Figure 50. IMS database recovery process

Information for a database recovery can come from any or all of the following sources:

- Image copies of the database
- Database reorganization data sets
- Log data sets (SLDSs and RLDs)
- Change accumulation data sets

You can use DBRC to track all of these information sources, greatly simplifying the task of database recovery.



**Related reading:** Refer to “Understanding the recovery task” on page 89 for more information about the recovery process.

If you register recoverable databases in the RECON data set, DBRC records the association of the databases to the log data sets containing database change records.

DBRC also records information about:

- Database image copies
- Reorganizations (except DEDB online reorganizations)
- Recoveries
- Change accumulations
- Backout

DBRC can generate JCL for executing a database recovery, because DBRC records this information in the RECON data set. Whether you use the GENJCL commands to generate JCL or provide the JCL yourself, DBRC uses information in the RECON data set to determine exactly which data sets are required for input. The Database Recovery utility runs only if DBRC verifies that the JCL is correct.

You can omit all logged changes after a certain time from the input by performing a *time-stamp recovery*. A time-stamp recovery is equivalent to backing out the omitted changes from the database.

Most time-stamp recoveries require DBRC in order to be successful. When you involve DBRC in your request for a time-stamp recovery, DBRC selects the correct logs and, at execution time, communicates to the Database Recovery utility where to stop processing the input to correctly perform your request.

The following figure shows how DBRC works with the Database Recovery utility.




Figure 51. How DBRC works with the Database Recovery utility

**Recommendation:** Implement DBRC in phases, defining at first only a few recoverable databases in the RECON data set. This allows you to gain experience in the use of DBRC, and gives you an opportunity to assess, make, and test any changes needed in your backup, recovery, and operational procedures.



---

## Chapter 32. How DBRC helps in recovery

DBRC is invoked by recovery tools and utilities to provide information about the resources required for database recovery. These resources can include information about image copies, logs, change accumulations, and database data sets.

---

### Generating recovery JCL

You can use either the GENJCL.RECOV command or the /RMGENJCL command to generate the JCL that is necessary to recover a registered database data set.

Using information recorded in the RECON data set, DBRC:

1. Selects the image copy data set to use for loading the most recent image copy
2. Selects the change accumulation and log data sets that are to be input to apply all the changes that were logged since the image copy was created

If change accumulation input is required (because of data sharing), but it is not present or usable, DBRC informs you of that fact and the GENJCL.RECOV command fails.

The GENJCL.USER command can generate user-defined output, which can include JCL. No skeletal JCL execution members are supplied to support the GENJCL.USER command. If you want to enter GENJCL.USER commands, you must supply the members to support them.

When you issue either the GENJCL command or the /RMGENJCL command, DBRC reads skeletal JCL and replaces symbolic parameters with actual values based on the information recorded in the RECON data set to build the appropriate JCL. For example, if you request that DBRC generate JCL to recover a database, DBRC retrieves the skeletal JCL member from the library and completes the JCL with information about the latest image copy, change accumulation, and log data sets, if necessary. Your databases must be registered in order for DBRC to generate JCL to process them.

The amount of time and effort required to recover a database can be significantly reduced by using the GENJCL command to generate the JCL and control statements necessary for the recovery. Using the GENJCL command also eliminates the causes of many recovery errors. You could spend a large amount of time during database recoveries determining which input data sets should be provided in what order to the Database Recovery utility.

When change accumulation data sets or PRILOG records (in the RECON data set) are available, DBRC selects them rather than the SLDS for recovery. This results in quicker database recoveries if you run the Database Change Accumulation regularly. DBRC knows which log data sets are required and ensures that IMS processes all volumes in the correct order. DBRC also selects the most recent image copy for database recovery.

DBRC always selects the optimum input for the Database Recovery utility by using change accumulation data sets whenever possible. If you have not used the Database Change Accumulation utility, or if that utility did not process some log

data sets, DBRC selects the required log data sets from the PRILOG (or SECLOG) records, which can contain RLDS, SLDS, or both RLDS and SLDS entries.

**Related reading:**

- See *IMS Version 11 System Definition* for more information about the tailoring actions for IMS PROCLIB members, the DBRC procedure, and the JCLOUT and JCLPDS DD statements.
- See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for details about customizing your own skeletal JCL and about the contents of IMS-supplied skeletal JCL.

**Recommendation:** For increased availability of data entry databases (DEDBs), use the DEDB Area Data Set Create utility to provide additional usable copies of an online area. It does not provide backup copies for recovery. The DEDB Area Data Set Create utility uses the RECON data set as part of its input.

---

## Validating utility JCL

When the a recovery utility processes a registered database data set, the utility presents its input to DBRC for validation. Whether the recovery JCL was created by you or by DBRC, DBRC verifies that the input JCL to the utility is correct (according to the current information in the RECON data set). It is possible, even if you created the JCL with the GENJCL command, that intervening events could invalidate the input JCL before the utility is run.

DBRC is invoked by the following IMS utilities and services to validate input and record the results:

- Index/ILDS Rebuild utility (DFSPREC0)
- Database Image Copy utility (DFSUDMP0)
- Database Image Copy 2 utility (DFSUDMT0)
- Online Database Image Copy utility (DFSUICP0)
- Database Change Accumulation utility (DFSUCUM0)
- Batch Backout utility (DFSBB000)
- Database Recovery utility (DFSURDB0)
- Log Recovery utility (DFSULTR0)
- Log Archive utility (DFSUARC0)
- HALDB online reorganization
- HD Reorganization Unload utility (DFSURGU0)
- HD Reorganization Reload utility (DFSURGL0)
- HISAM Reorganization Unload utility (DFSURUL0)
- HISAM Reorganization Reload utility (DFSURRL0)
- Database Prefix Update utility (DFSURGP0)
- DEDB Area Data Set Create utility (DBFUMRI0)
- /RECOVER commands

The following figure shows DBRC's role in running the previously mentioned utilities.




Figure 52. DBRC's role in utility execution

**Related reading:** For more information about:

- The IMS database utilities, see *IMS Version 11 Database Utilities*.
- The IMS log utilities, see *IMS Version 11 System Utilities*.
- The NOTIFY.BKOUT command, which manually creates a BACKOUT record in the RECON data set, see *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*.
- The /RECOVER command, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

When you run the Batch Backout utility (DFSBB000), DBRC determines the complete set of logs that are needed for a particular backout job. In addition, DBRC manages information about the logs so that backout and restart jobs can be easily coordinated.

**Exception:** DBRC does not verify the JCL input for the HD and the HISAM Reorganization utilities, but does record information about their execution in the RECON data set.



---

## Chapter 33. Recording information about opening and updating databases

After IMS opens a database, DBRC passes back the RECON data set initialization token (RIT) and any extended error queue elements (EEQEs) associated with each DBDS. The RIT allows IMS to determine whether the database has been used without DBRC or whether the database has been controlled by a different RECON data set.

**Related reading:** See *IMS Version 11 Database Administration* for information on EEQEs.

When changes to DBDSs and areas occur, DBRC records information about these changes in the RECON data set. DBRC subsequently uses this information to determine which log data sets might contain change records for a given DBDS or area.

When a DBDS that is registered in the RECON data set is first opened for updates (or allocated), IMS tells DBRC to create an ALLOC record. In the case of a DEDB area, the ALLOC record is created when the area is first opened for update. The ALLOC record identifies the DBDS or area and contains the time stamp of the first update and the open time stamp of the corresponding PRILOG.

When DBRC creates the ALLOC record, DBRC also enters the name of the DBDS or area being changed in the LOGALL record for the PRILOG that is active at the time of the change.

When you deallocate (close) a DBDS or area using a /DBRECOVERY command from the operator console of the online IMS subsystem, DBRC writes a deallocation time stamp in the ALLOC record. If no de-allocation time is recorded, DBRC uses the closing time of the associated log as the deallocation time. Thus the RECON data set contains a list of the names of DBDSs or areas for which change records might exist on a given log data set (LOGALL record) and a list of the time ranges where changes could exist for a specific DBDS or area (ALLOC records) and a list of the logs containing the changes.





---

## Chapter 34. Supporting data sharing

Data sharing requires that the databases be registered with DBRC. DBRC checks that subsystems have authority to perform the requested task and that other subsystems are not currently reserving the database.

**Related reading:** See Chapter 15, “Data sharing in IMS environments,” on page 233 and *IMS Version 11 Operations and Automation* for more information on data sharing.

---

### Levels of data sharing

DBRC supports the two levels of IMS data sharing: database level and block level.

#### Database level

The entire database or DEDB area is a resource that can be accessed for update by a single IMS system at a time. For area resources this can also be called *Area-level sharing*.

#### Block level

A database or DEDB area can be accessed by multiple IMS subsystems concurrently. Data integrity is preserved for the IMS subsystems that access the shared data. Within a database or area, resources are reserved at the block level.

**Definition:** For OSAM databases, the block is a physical data block stored on DASD. For VSAM databases and DEDBs, the block is a control interval (CI).

---

### Recording data-sharing information in the RECON data sets

DBRC records data-sharing information in the RECON data set for all registered databases.

The following data-sharing information is recorded into the RECON data set:

- Sharing level allowed for each database
- Names of databases or areas currently authorized for processing
- Names of IMS subsystems that are involved
- Statuses of the IMS subsystems
- Database statuses from a recovery viewpoint

---

### Assigning a sharing level with DBRC

The sharing level of a database or DEDB area determines whether a request for access is granted. DBRC allows you to establish one of four sharing levels.

The following sharing levels are defined using the INIT.DB command and modified with the CHANGE.DB command.

#### SHARELVL 0

The database is not to be shared. The database can be authorized for use by one IMS system at a time. SHARELVL 0 is equivalent to specifying ACCESS=EX on the /START command.

**SHAREVL 1**

Sharing is at the database level. One IMS system can be authorized for update at one time; any sharing systems can only be authorized for read-only processing. Otherwise, the data sharing is for multiple readers.

**SHAREVL 2**

Sharing is at the block level but only within the scope of a single IRLM and a single z/OS. Sharing requires that IMS subsystems sharing a database use the same RECON data set. Multiple IMS systems can be authorized for update or read processing.

**SHAREVL 3**

Sharing is at the block level by multiple IMS subsystems on multiple Realms. Multiple IMS systems can be authorized for nonexclusive access. The IMS subsystems can be on multiple z/OS images using different IRLMs.

**Note:**

To ensure the integrity of databases in data sharing environments when batch jobs running without IRLM support access SHAREVL 3 databases, DBRC authorizes batch jobs that have update access only if all other IMS systems and batch jobs that are currently authorized by DBRC to the database have read-only access.

If IRLM=N is specified in the IMSCTRL macro or DBBATCH procedure, DBRC authorizes a batch job for update access to a database only if all other IMS systems and batch jobs that are currently authorized by DBRC to the database have read-only access.

When a batch IRLM=N job has authorization for update, authorization for an online system or other batch job fails unless it is for read-only access.

---

## Chapter 35. DBRC support for Remote Site Recovery

DBRC assists you with the definition and management of IMS components in the RSR complex.

In support of RSR, DBRC provides:

- Commands to define, update, and display the status of the RSR complex.

The RECON data set contains the definition of an RSR complex.

- Services that are used by an active subsystem to identify the tracking subsystem and the databases covered by RSR.

An active subsystem obtains the identity of its tracking subsystem from DBRC.

As databases are updated by the active subsystem, DBRC tells the database component whether the database is covered by RSR. If the databases are being tracked by RSR, the active subsystem sends its log data to the tracking subsystem.

- Services used by a tracking subsystem to record information about log data that is received from an active subsystem.

As logs are received and stored at the tracking site, DBRC records the receipt of the log data. When begin-update records are received for registered databases, DBRC records the database update.

- Tracking subsystem database support:

- Two types of tracking (called shadowing): DB level tracking (DBTRACK) or Recovery level tracking (RCVTRACK).
- Maintains log data set information for online forward recovery.
- Records which database change records have actually been applied to the covered databases.

- Services to assist in the takeover process

During a remote takeover, DBRC changes the state information of the registered databases at the new active site to indicate that they are now the active databases.

### **Related reading:**

- See Chapter 45, “Remote Site Recovery overview,” on page 663 for more information about RSR.
- See “Recovery in an RSR environment” in *IMS Version 11 Database Administration* for more information on controlling database recovery in an RSR environment.



---

## Chapter 36. Supporting IMSplexes

One DBRC function that is available only in an IMSplex is that when an I/O error occurs on a RECON data set in an IMSplex and a spare data set is available, the instance of DBRC that noticed the error copies the good RECON data set to the spare, activates the spare, and deallocates the original RECON data set.

At this point in the processing, the DBRC that noticed the I/O error can automatically notify the other DBRCs in the IMSplex about the reconfiguration. Then, after the original RECON data set is deallocated, it can be deleted and redefined as the spare RECON data set.

You can access the RECON data set in either serial or parallel mode, where all the IMS systems are at Version 10 or later. In serial mode, the RECON data set is available only to one DBRC instance at a time because that DBRC uses a z/OS RESERVE macro to lock the entire RECON data set while the update is in progress. In parallel mode, DBRC uses the Transactional VSAM function of DFSMS to provide record-level sharing for the RECON data set. All access to the RECON for a single DBRC request (from IMS or a command or API) is processed by Transactional VSAM as one or more transactions.

**Related reading:** For more information about accessing the RECON data set in parallel mode, see “Accessing the RECON data sets in parallel mode” on page 536.



---

## Chapter 37. Defining DBRC to IMS

You define DBRC to IMS as part of the IMS system definition process. This topic provides a brief overview of the requirements for creating and executing a DBRC region.

### **Related reading:**

- See Chapter 41, “Initializing and maintaining the RECON data sets,” on page 535 for information on creating and allocating the RECON data set.
- See Chapter 38, “Registering databases and database data sets,” on page 507 for information about registering databases.
- Refer to *IMS Version 11 Installation*, *IMS Version 11 Release Planning*, and *IMS Version 11 System Definition* for:
  - A complete description of the IMS installation procedures and requirements.
  - Information about the DBRC procedure and its parameters.
  - Information about the DFSIDEF0 module.
- See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for detailed information about the INIT.RECON command.

IMS online systems always use DBRC; you cannot override this.

You can choose whether IMS batch jobs can use DBRC or not, but certain functions, such as data sharing, cannot be used without DBRC.

Batch and utility regions can use the DFSIDEF0 module during initialization. In DFSIDEF0, you can set the DBRC= keyword to YES, NO, or FORCE. The DFSIDEF0 module that is shipped in the ADFSSMPL library contains a batch and utility region default of DBRC=YES. This value is coded on the DFSIDEF macro.

1. Use the DBRCNM= parameter to request that IMS create a cataloged DBRC procedure.

You can override the DBRC procedure name specified in your system definition in the DFSPBIMS, DFSPBDBC, and DFSPBDCC members.
2. Determine whether DBRC is used in a batch procedure.

If DBRC is going to be used in a batch procedure, specify this in the DBRC= EXEC parameter. This EXEC parameter is ignored by an online IMS.

The DBRCNM= parameter can be used to override the DBRC procedure name for an online IMS execution.
3. Place DBRC's load modules into a load library that is in the normal load library search sequence for your IMS load modules, for example, IMS.SDFSRESL.
4. Include the DBRC procedure during system definition.

IMS automatically starts the DBRC procedure by issuing a z/OS START command during control region initialization. This procedure specifies parameters for the DBRC region. DBRC runs in its own address space for an online IMS subsystem.

To include the DBRC procedure during system definition, copy the skeletal DBRC procedure from IMS.PROCLIB to SYS1.PROCLIB. The member name must match the name specified on the DBRCNM parameter in the IMSCTRL macro or the applicable EXEC procedure.



If DBRCNM is specified in more than one place, or if DBRCNM is not explicitly specified, the following order of precedence applies:

- DBRCNM=DBRC is the default.
- DBRCNM=*name* in the IMSCTRL macro overrides the default.
- DBRCNM=*name* defined in a DFSPBxxx member in IMS.PROCLIB overrides the IMSCTRL macro setting.
- DBRCNM=*name* defined in a JCL EXEC parameter overrides the IMS.PROCLIB member setting.

5. Initialize the RECON data sets.

Use the Access Method Services (IDCAMS) DEFINE CLUSTER command to create the RECON data sets and then use the INIT.RECON command to initialize the RECON data sets as usable by DBRC.

If you do not intend to register databases, the INIT.RECON command is the only command you need to issue in order to initialize the data set.

---

## Chapter 38. Registering databases and database data sets

The RECON data set must have a DB record for each database whose recovery DBRC is to control.

If you want DBRC to control database recovery, you must register the databases in the RECON data set. DBRC then records information about database updates and about the corresponding log data sets that contain updated log records. DBRC also records the creation of image copy and change accumulation data sets, and records database recoveries and reorganizations that affect registered databases.

For non-HALDBs (High Availability Large Databases), use the INIT.DB and INIT.DBDS commands to register databases in the RECON data set and to define them as recoverable or nonrecoverable. For HALDBs, you can use the INIT.DB and INIT.PART commands, or the HALDB Partition Definition utility.

For each non-HALDB database that you have registered, issue the INIT.DBDS command to register all its data sets or DEDB areas. For DEDBs, use the INIT.ADS command to identify the data sets within each area. An area can have up to seven area data sets (ADSs).

**Requirement:** Utility jobs that issue either an INIT.DB or INIT.PART command (to register a HALDB) or any INIT.DBDS command, require a DD name for the IMS.DBDLIB data set that contains an entry for the HALDB DBD (or non-HALDB DBDS) for which you are issuing the command. A DD name for the IMS.DBDLIB is not needed for non-HALDB command.

From a DBRC perspective, a HALDB consists of a HALDB master (TYPE=HALDB) and HALDB partitions (TYPE=PART).

To update or delete information about HALDBs in the RECON data set, you can use the HALDB Partition Definition utility or the following DBRC commands; CHANGE.DB, CHANGE.DBDS, CHANGE.PART, DELETE.PART, or DELETE.DB.

### **Related reading:**

- See *IMS Version 11 Database Administration* for an overview of HALDBs and detailed information about how to create them.
- See *IMS Version 11 Database Utilities* for information about the HALDB Partition Definition utility.
- See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for information about the CHANGE commands and specifics about the INIT commands.



---

## Chapter 39. Planning for recovery

Planning for recovery (in non-data sharing and data sharing environments) involves two primary tasks: setting up the mechanisms of recovery (logging, taking checkpoints, keeping records) and setting up your operational procedures for situations that require recovery.

---

### Setting up recovery mechanisms

The logging and checkpoint mechanisms of online IMS subsystems in a non-sharing environment are also active in a data sharing environment.

These include:

- System log data sets and the use of the WADS and restart data sets
- Program, system, and message-queue checkpoints
- Making image copies of databases

The primary difference between non-sharing and data sharing environments is in their degree of reliance on DBRC. DBRC helps control the data-sharing environment; it does not merely keep records.

Specifically, the following items are different in a non-data sharing environment:

- Database reconstruction (forward recovery)
- Program-level (dynamic) backout
- Database backout
- IMS restart (with embedded recovery)

**Related reading:** See *IMS Version 11 Operations and Automation* for detailed information on recovery procedures in non-data sharing and data-sharing environments.

---

### Recovery facilities

Dynamic backout and batch backout support behave differently in non-data sharing and data-sharing environments.

- Dynamic backout

In non-data sharing and data-sharing environments, an online IMS subsystem dynamically backs out the uncommitted database changes of an application program and discards its uncommitted output messages under either of these conditions:

- The program fails.
- The program requests backout with a rollback call.

In a data-sharing environment, however, IRLM locks and status indicators in the RECON data set ensure integrity by protecting uncommitted changes from sharing subsystems. Operation of the system and other programs continues uninterrupted.

**Related reading:** See *IMS Version 11 Database Administration* for detailed information about dynamic backout.

- Batch backout support

DBRC improves database integrity by interfacing with IMS restart dynamic backout and the Batch Backout utility to control access to databases in need of backout. DBRC creates a BACKOUT record to track each unit-of-recovery (UOR) for each database in need of backout. DBRC verifies logs that are input to the Batch Backout utility.

BACKOUT records are created for online subsystems. BACKOUT records are not created for DL/I batch subsystems unless dynamic backout was being used and it failed.

Prior to performing an emergency restart with the COLDSYS parameter (/ERESTART COLDSYS), run batch backout with either the COLDSTART or ACTIVE parameter. After the batch backout completes, DBRC creates a BACKOUT record for all in-flight and indoubt UORs as candidates for backout. The next IMS restart promotes candidate UORs to backout-needed status and DBRC sets BACKOUT NEEDED=ON in the database records in the RECON data set.

When dynamic backout fails, a BACKOUT record is created with a UOR indicating dynamic backout failure and DBRC sets BACKOUT NEEDED=ON in the database record in the RECON data set.

When the databases are backed out successfully, DBRC resets the backout flags in the database records (in the RECON data set) appropriately and updates the BACKOUT records. When backout processing for all of the UORs have been completed, DBRC deletes the BACKOUT record from the RECON data set.

For DBCTL, if you choose not to back out an unresolved indoubt UOR, use either the DELETE.BKOUT or CHANGE.BKOUT command to remove the unresolved indoubt UOR from the BACKOUT record in the RECON data set.

DBRC commands are available to update BACKOUT records if needed. The need to make manual changes to the backout record should be minimal.

**Attention:** Use the DELETE.BKOUT or CHANGE.BKOUT commands with extreme caution because you could compromise database integrity.

DBRC verifies the validity of the logs used as input to the Batch Backout utility. DBRC checks the input log volumes to ensure that they are in the correct sequence, that all the needed logs are provided, and that they are properly closed. When you include ACTIVE or COLDSTART statements in the Batch Backout utility SYSIN statement, DBRC performs an additional check to ensure all volumes related to restart are included. For DL/I batch logs, a check is also performed to ensure that the correct volumes are from the last batch execution.

**Attention:** If BYPASS LOGVER is included in the SYSIN statement of the Batch Backout utility, then DBRC does not verify the input log and the Batch Backout utility does not notify DBRC about the UORs. Existing UORs are updated when backout processing has completed successfully.

Although DBRC's backout support protects databases from damage caused by the most common errors, the following list describes some other possible sources of damage to databases:

- If an IMS subsystem abends and an /ERE COLDSYS command is issued before the in-flight and indoubt UORs have been identified to DBRC (by a Batch Backout run), the databases associated with those in-flights and in-doubts are not protected from erroneous updates until the first Batch Backout run using the COLDSTART or ACTIVE statements.
- Including logs from multiple runs of a batch job in the same log data set (by specifying DISP=MOD) makes log verification unreliable.
- You can modify the RECON data set with DBRC commands. Careless use of these commands could lead to errors, such as:

- Allowing backouts that should not be performed.
- Allowing other subsystems access to databases that need to be backed out.
- The Batch Backout BYPASS LOGVER control statement allows backouts to be performed when information in the RECON data set indicates that the input log is invalid or that the backouts are not needed. Careless use of this control statement can cause backouts to be performed that should not be performed.
- Unregistered databases are not protected from being used while backout is needed.
- Backing out a normally terminated job (that did not use IRLM) after a time-stamp recovery was performed (to recover the database to a point in time prior to this log) makes log verification unreliable. DBRC will erroneously validate that the log can be used as input because the log being backed out is the last log for this subsystem name (SSID).

**Related reading:** See *IMS Version 11 Database Administration* and *IMS Version 11 Database Utilities* for more information about UORs, the database backout process, and the Batch Backout utility.

- Forward recovery

The process of recovering a database in a data-sharing environment has certain similarities to recovering a database in a non-data sharing environment. In both environments, use DBRC for database recovery and use the following as input to the recovery process:

- The most recent image copy of the lost database

**Note:** If HALDB Online Reorganization is used for the starting point of recovery, then an image copy is not needed because all database changes will be on the log data set(s).

- The recovery log data set (RLDS) or all pertinent system log data sets used since that copy was made

**Recommendation:** Use the RLDS as input to the recovery process. The RLDS is a more efficient input to the recovery process than the pertinent system log data sets because the RLDS contains only database change log records.

If you are using the Database Recovery Control utility for forward recovery and block-level data sharing is involved, you might require the additional step of change accumulation.

**Related reading:** See *IMS Version 11 Database Administration* for further discussion of forward recovery.

---

## Recovery without DBRC

If you perform any recovery-related actions offline when DBRC is not running, such as making database image copies, problems could arise because DBRC would not be notified of the changes in status.

Therefore, DBRC must be specifically informed of such changes. DBRC offers you several commands for this purpose.

Because restart is required, the Utility Control Facility (UCF) might have to be used without DBRC being active. If you use UCF for any recovery-related work, DBRC must be informed of status changes afterward.

**Related reading:** See *IMS Version 11 Database Administration* for detailed information about recovery without DBRC and its related commands.

---

## Restart after an IMS failure

Restart an IMS subsystem in a data sharing environment in the same way as in a non-data sharing environment.

- If you can shut IMS down normally (a checkpoint shutdown using the /CHECKPOINT PURGE, /CHECKPOINT DUMPQ, or /CHECKPOINT FREEZE commands), use the /NRESTART command to restart IMS normally.
- If the IMS failed, you must use the /ERESTART command to perform an emergency restart.

In a data-sharing environment, if the associated IRLM also stops or fails, you must first start IRLM before you start IMS in a block-level environment.

**Related reading:** See *IMS Version 11 Operations and Automation* for more information about restarting IMS after it fails.

---

## Restart after a DBRC failure

IMS knows whether its DBRC fails because DBRC runs under the control of IMS. If a DBRC failure occurs, IMS terminates abnormally, attempts to flush its buffers, and closes the system log. After correcting the DBRC problem, you can restart IMS by issuing the /ERESTART command.

---

## Recovery involving IRLM configurations

DBRC records IRLM status information in the subsystem record in the RECON data set.

**Related reference:**

 Fields in a RECON listing, by record type (Commands)

---

## Batch backout

Prior to executing an emergency restart with the COLDSYS parameter (/ERESTART COLDSYS), run the Batch Backout utility with either the COLDSTART or ACTIVE parameter.

Provide DBRC with the names of the logs to protect registered databases (that need to be backed-out) from being accessed until their backouts are complete.

For DB level control, if you choose not to back out an unresolved in-doubt UOR, use the DELETE.BKOUT command to remove it from the BACKOUT record.

**Attention:** Use the DELETE.BKOUT command with extreme caution. It deletes all backout information for a subsystem from the RECON data set.

---

## Chapter 40. Considerations for a DBRC system

DBRC controls database-related and log-related processes.

These processes include:

- Creating backup copies of databases
- Creating DB change accumulation data sets
- Recovering databases
- Protecting databases that need backout
- Archiving OLDSs

---

### General considerations for using DBRC

DBRC has several considerations to be aware of during its use.

Be aware of the following considerations when using DBRC:

- DBRC does not support main storage databases (MSDBs).
- DBRC plays no role in the processing of GSAM databases, so there is no reason to register them.
- Logging is required for batch jobs that use DBRC and have update access.
- Never update the RECON data set information (such as for DBDSs or log data sets) about a data set that is currently in use.
- Be sure to set the time-of-day (TOD) clock value in the host processors accurately, or unpredictable results might occur in the RECON data set.

---

### Data set naming conventions

In most cases, the use of z/OS catalog facilities for image copy and change accumulation data sets is optional because DBRC always records volume serial information pertaining to these data sets in the RECON data set. If you catalog image copy and change accumulation data sets, they must have unique data set names.

**Requirement:** Image copies taken using the DFSMSdss fast replication copy function of the Database Image Copy 2 utility must use z/OS catalog facilities.

DBRC provides a data set naming convention to help you generate unique data set names for those image copy data sets (for HSSP image copies and concurrent image copies, as well as standard image copies) and change accumulation data sets that you define for future use.

If you use this convention all the time, uniqueness of your data set names is assured. If you use the convention only occasionally, you are sent a message at the end of your job step that indicates that you did not follow the naming convention and that duplicate data set names could exist in the RECON data set. DBRC assumes that data set names specified in quotation marks do not follow the naming convention. Therefore, DBRC does not check data set names surrounded by quotation marks.



When you add records to the RECON data set that create data sets for one of the recovery utilities to use in the future and you are using this data set naming convention, you can specify either the fully-qualified data set name or simply the abbreviation *high-level-qualifier.\*.low-level-qualifier* (for example, ALPHA1.\*.OMEGA). You can use these abbreviated names on any INIT, CHANGE, DELETE, or NOTIFY.REORG command of DBRC when you are specifying the name of a data set that follows the naming convention. DBRC expands the abbreviated name to its fully-qualified form before it accesses the RECON data set.

The following topics describe additional data set naming conventions.

## Naming convention for image copy data sets

The naming convention for image copy data sets that do not use the DFSMSdss fast replication function is *high-level-qualifier.dbdname.ddname.IC.low-level-qualifier*.

### *high-level-qualifier*

A character string that can be from 1 to 8 alphanumeric characters long. The first character must be alphabetic.

### *dbdname*

The database name of the DBDS for which the image copy data set is being recorded in the RECON data set.

### *ddname*

The data set DD name of the DBDS for which the image copy data set is being recorded in the RECON data set.

**IC** Indicates that this is the image copy data set.

### *low-level-qualifier*

A character string that can be from 1 to 8 alphanumeric characters long. It must be unique for each DBDS and the first character must be alphabetic.

The naming convention for image copy data sets that use the DFSMSdss fast replication function can be either of the two following formats:

*high-level-qualifier.dbdname.ddname.Dyyddd.Thhmmss*

*high-level-qualifier.dbdname.ddname.*

### *high-level-qualifier*

A character string that can be from 1 to 25 alphanumeric characters long. The first character must be alphabetic. The character string may represent multiple data set qualifiers, each from 1 to 8 alphanumeric characters long, separated by a period. The first character of each qualifier must be alphabetic, and the last character of the string must not be a period.

### *dbdname*

The DBD name of the database for which the image copy data set is being recorded in the RECON data set.

### *ddname*

The data set DD name of the DBDS for which the image copy data set is being recorded in the RECON data set.

*yyyddd* Is the date. This string must begin with character string 'D'.

### *hhmmss*

Is the time stamp. This string must begin with character string 'T'.

**Related reading:** For the details of using the DFSMSdss fast replication function with the Database Image Copy 2 utility (DFSUDMT0), see the information about this utility in *IMS Version 11 Database Utilities*.

## Naming convention for duplicate image copy data sets

The format for duplicate image copy data sets is *high-level-qualifier.dbdname.ddname.IC2.low-level-qualifier*.

This format is identical to the convention for image copy data sets, except that the IC2 field indicates that this is a duplicate image copy data set.

## Naming convention for change accumulation data sets

The format for change accumulation data sets is: *high-level-qualifier.cagrpname.CA.low-level-qualifier*.

*high-level-qualifier*

A character string that can be from 1 to 8 alphanumeric characters long. The first character must be alphabetic.

*cagrpname*

The name of the CA group for which you are creating the change accumulation data set.

**CA** Indicates that this is a change accumulation data set.

*low-level-qualifier*

A character string that can be from 1 to 8 alphanumeric characters long and must be unique for each CA group. The first character must be alphabetic.

---

## Database backup copies

When IMS takes a regular system checkpoint, it records internal control information for DL/I (for Fast Path, IMS also records buffers and MSDBs), but it does not record the external contents of the database. If the database is lost, examining the last system checkpoint does not help. The log can tell you what changes have occurred, but without a backup copy of the database, recovery is impossible.

**Recommendation:** Make a backup copy of the database after you initially load it, and make new backup copies at regular intervals. The more recent the backup copy is, the fewer log change records need to be processed during recovery, thus reducing the time that is needed for recovery.

IMS enables you to make backup copies in these ways:

- The Database Image Copy utility (DFSUDMP0) runs offline and uses access method services to make the copy.
- The Database Image Copy 2 utility (DFSUDMT0) runs offline and invokes either DFSMSdss Concurrent Copy or DFSMSdss fast replication to make the copy.
- The Online Database Image Copy utility (DFSUICP0) runs online and uses IMS services to make the copy.
- The unloaded data that is output from the HISAM Reorganization Unload utility (DFSURUL0) can be used as a backup copy.
- HSSP processing can also create image copies of DEDB areas.

When these utilities run, they can (depending on installation parameters) call DBRC to update essential information in the RECON data set. If you use a supported image copy utility, DBRC records the image copies for registered databases. DBRC also generates the JCL for the utility if you enter the GENJCL.IC or GENJCL.OIC command.

You can also use various utilities supplied by the operating system to make your backup copies; however, these do not interact with DBRC, and so you need to take certain actions to notify DBRC of your non-standard image copies.

**Related reading:** See *IMS Version 11 Database Administration* for more information on HSSP processing of DEDB areas.

## Image copy utilities (DFSUDMP0, DFSUDMT0, DFSUICP0)

IMS enables you to “take a picture” of your database before and after changes have been made to the database. The “pictures” are called *image copies*. The term refers to the fact that the copy is an as-is image; the image copy utilities do not alter the physical format of the database as they copy it. Image copies are backup copies of your data that help speed up the process of database recovery.

The Database Image Copy utility (DFSUDMP0), Database Image Copy 2 utility (DFSUDMT0), and Online Database Image Copy utility (DFSUICP0) create image copies of databases. All of the image copy utilities operate on data sets or DEDB areas, so if a database is composed of multiple data sets or areas, be sure to supply the utility with multiple specifications. You can request that one of the supported image copy utilities produce both an image copy data set and a duplicate image copy data set in one run of the utility.

**Recommendation:** Copy all data sets or areas belonging to a database at one time. If you perform multiple recoveries in order to reset a database to a prior state, recover all data sets belonging to the database and to all logically related databases (including those related by application processing) to the same point to avoid integrity problems.

Each of the image copy utilities provide the option to create backup copies without taking databases and areas offline. You can use this capability to provide increased database availability. Image copies taken while the database is available for concurrent update processing by IMS applications are called concurrent image copies or *fuzzy* image copies. Changes already made to the database by active applications might be missing from the copy because the changes might not have been physically written to the data set. These changes, however, have been written to the log. In this case, it is necessary to go back to some earlier point in the logs to ensure that all changes are applied. How far to go back depends on the type of database and which image copy utility was used.

When concurrent copy or fast replication image copies are not used, the database must be either taken offline or made available only for 'read' access and a consistent or 'clean' image copy is taken. See “Concurrent image copy” on page 519 for more information.

If the image copy was made while the database was not being accessed for update, only changes that were logged after the run time of the copy are required.

When using these utilities, you have the option of creating one to four output image copies. Only the Database Image Copy 2 utility allows three or four output

copies and only the first two output copies are recorded in the RECON data set. When the Database Image Copy 2 utility uses the DFSMSdss fast replication function, only one output copy is created.

The advantage of making multiple copies is that if an I/O error occurs on one copy, the utility continues to completion on the other copies. Also, if one copy cannot be read, you can perform recovery using another. The trade-off in deciding whether to make multiple copies, is that performance can be degraded because of the time required to write the additional copies.

DBRC works similarly with the three image copy utilities. When the DFSMSdss fast replication function is not used, the rules for pre-definition and reuse of image copy data sets apply to all three. Each utility calls DBRC:

- To verify the input to the utility (DBRC allows it to run only if the input is valid.
- To record information in the RECON data set about the image copy data sets that it creates.

An image copy record in the RECON data set has the same format whether its corresponding image copy data set was created by the Database Image Copy utility, the Database Image Copy 2 utility, or by the Online Database Image Copy utility.

Two different commands create image copy jobs:

- GENJCL.IC for the offline utilities Database Image Copy and Database Image Copy 2
- GENJCL.OIC for the online utility Online Database Image Copy

When you run batch jobs without logging, take an image copy immediately afterwards; do not count on rerunning the batch jobs, as part of a subsequent recovery, in combination with the Database Recovery utility. The database could be damaged by the combination because the batch processing is not guaranteed to be physically repeatable.

### **Database image copy (DFSUDMP0)**

The Database Image Copy runs offline. It supports a concurrent image copy (CIC) option that enables you to create an image copy while the database remains online. The CIC option can create image copies of OSAM data sets and VSAM Entry Sequenced Data Sets (ESDSs). The Database Image Copy does not support creating image copies of VSAM Key Sequenced Data Set (KSDSs).

Use the Database Image Copy (DFSUDMP0) utility to create data set copies of:

- HISAM databases
- HIDAM databases
- HDAM databases
- PHDAM (partitioned HDAM databases)
- PHIDAM (partitioned HIDAM) databases
- PSINDEX (partitioned secondary index) databases
- DEDB areas

To request a concurrent image copy, use the CIC keyword on the GENJCL.IC command. Alternatively, you can specify the CIC parameter on the EXEC

statement for image copy job. DBRC must be used by the utility and you can only take a concurrent image copy of a database that has been registered with DBRC.

When you run the Database Image Copy utility to take a consistent image copy (CIC option not specified), DBRC is not required but is recommended. DBRC ensures that there is no update activity against the database or area while the utility is executing. If you run the utility without using DBRC you must make certain that no updates occur to the database or area. You can issue a /DBDUMP command or a /STOP AREA command, for example, to prevent updating of the database or area by transactions in the system previously doing updates.

### **Database image copy 2 (DFSUDMT0)**

The Database Image Copy 2 utility (DFSUDMT0) uses either the DFSMSdss concurrent copy function or the DFSMSdss fast replication function to take an image copy.

When using the concurrent copy option, DFSUDMT0 invokes DFSMSdss to dump the input data set. The output is recorded in the RECON data set and can be used for database recovery in the same way as the output from either of the other image copy utilities. You can use this utility for HISAM, HIDAM, HDAM, PHDAM, PHIDAM, and PSINDEX databases and for DEDB areas. Database data sets that are to be copied by this utility must reside on hardware that supports the Concurrent Copy option.

An optional COMPRESS parameter invokes the “compress” option in DFSMSdss concurrent copy. The compress option enables you to reduce the storage space required to hold the image copy; however, using the compress option increases the CPU time that is required to perform the copy operation.

**Note:** The COMPRESS parameter is not available to DFSMSdss fast replication.

When using the fast replication option, DFSUDMT0 invokes DFSMSdss to COPY the input data set. The output is recorded in the RECON data set and can be used for database recovery in the same way as the output from either of the other image copy utilities. You can use this utility for HISAM, HIDAM, HDAM, PHDAM, PHIDAM, and PSINDEX databases and for DEDB areas. Database data sets that are to be copied by this utility must reside on hardware that supports the fast replication option.

The Database Image Copy 2 utility provides greater database availability when taking consistent image copies. The database needs to be unavailable for update processing for only a very short period of time while DFSMSdss establishes a concurrent copy or fast replication session. Update processing can then be resumed while the image copy data sets are actually being written. The updates are not included in the image copy.

This utility can also copy the database while it is being updated by IMS applications. The image copy created in this case is a “fuzzy” image copy. The Database Image Copy 2 utility can copy all supported data set types, including VSAM KSDSs, while the databases remain online. In order to copy a VSAM KSDS while it is online, it must be SMS managed and allocated with IDCAMS using the BWO(TYPEIMS) parameter.

**Requirement:** The Database Image Copy 2 utility must use DBRC and the databases and areas being copied must be registered with DBRC. The utility can

create up to 4 copies using concurrent copy; only 1 copy can be made using fast replication. Only the information about the primary and secondary copies are recorded in the RECON data set.

The concurrent copy image copy output from the Database Image Copy 2 utility is in DFSMSdss dump format, which is different than the format of the output of the other image copy utilities. It is, however, directly usable as input to database recovery. Information about the image copies are recorded in the RECON data set with image copy type SMSCIC, or SMSNOCIC.

The image copy output from a fast replication invocation of the Database Image Copy 2 utility is in DFSMSdss copy format, which is different than the format of the output of the other image copy utilities. Like the DFSMSdss dump format, it is directly usable as input to database recovery. Information about the image copies are recorded in the RECON data set with image copy type SMSONLC or SMSOFFLC.

**Requirement:** The Database Recovery utility must use DBRC when using the SMSCIC, SMSNOCIC, SMSONLC, or SMSOFFLC image copy data sets for recovery.

### **Online database image copy (DFSUICP0)**

Online Database Image Copy (DFSUICP0) runs as a BMP program in the online IMS and DBCTL environments.

You can use it for HISAM, HIDAM, HDAM, PHDAM, PHIDAM, and PSINDEX databases only; it does not support areas. A database being copied by this utility is available for updating by the IMS subsystem in which the utility is executing. Other IMS subsystems cannot have concurrent update access to the database.

If the database being copied is updated while the utility is running, a fuzzy image copy is produced. Recovery with this image copy requires all logs starting with the log that was in use when the Online Database Image Copy utility was started.

## **Concurrent image copy**

IMS provides the capability to take an image copy of a database without taking that database offline. This means the database can be updated while the image copy is being taken and some, all, or none of the updates might appear in the image copy.

This image copy is called “fuzzy” because the copy represents the state of the database over a certain period of time rather than at one specific instant in time. It is also called a 'concurrent image copy' because the copy was taken while update processing is happening.

The ability to take image copies while the databases are being updated by IMS applications allows increased database availability. The offline image copy utilities, Database Image Copy and Database Image Copy 2, provide an option to take a concurrent image copy. Concurrent image copies can be created by both the concurrent copy and fast replication options of Database Image Copy 2. A database being copied by the Online Database Image Copy utility can be concurrently updated by the IMS subsystem in which the utility is running (but not by other IMS subsystems). Image copies created by HSSP processing are also 'fuzzy' copies because the areas are available for update processing while HSSP is running.



When a consistent image copy is input to database recovery, the recovery only requires logs from after the image copy job completed. A concurrent image copy, might not include updates that were made before the image copy process started or while it was executing. Therefore, when a concurrent image copy is input to recovery, logs from before the image copy process was started might have to be supplied to database recovery.

**Restrictions:** The following restrictions apply to creating concurrent image copies:

- Only databases that are registered with DBRC are eligible for concurrent image copying.
- The Database Image Copy utility (which can take consistent image copies without using DBRC) must use DBRC when creating a concurrent image copy.
- Nonrecoverable databases are not eligible for concurrent image copying because there is no log data to complete the image copy.

**Related concepts:**

“Database backup copies” on page 111

## Creating image copy data sets for future use and reuse

Use the REUSE parameter to inform DBRC that you want to be able to define image copy data sets and record them in the RECON data set for future use.

You define image copy data sets with the INIT.IC command. When processing the GENJCL.IC command, DBRC selects one of the image copy data sets for use by the image copy utility.

**Restriction:** DFSMSdss fast replication, as invoked by the Database Image Copy 2 utility, does not support the REUSE parameter.

When the Database Image Copy utility uses an available image copy data set, DBRC updates its record in the RECON data set with the time stamp of the run of the Database Image Copy utility during which the image copy data set was used.

If you specify the NOREUSE parameter, you cannot predefine image copy data sets (This is the default). You need to supply the output data set name for the utility in either the skeletal JCL member used in processing the GENJCL.IC command or in the JCL that you produce yourself. When you specify NOREUSE, DBRC dynamically sets the unit type of the output image copy data set. DBRC sets it to the default unit type for the device as specified in the INIT.RECON and CHANGE.RECON commands. Specify NOREUSE when you want more than two DFSMSdss concurrent copies (you can have up to four).

If you do not specify REUSE, every time the image copy utility is run DBRC deletes the oldest image record that exceeds both the GENMAX and RECOVPD values. The image copy data set itself is not deleted-only its record in the RECON data set is deleted. You must either delete the data set or keep track of it yourself, because DBRC is no longer aware of its existence.

If you are using the image copy option of HSSP for a DEDB area, the area must be defined with the REUSE parameter and the data sets you predefine must be cataloged.

## Controlling the number of image copies managed

Use the INIT.DBDS and the CHANGE.DBDS commands to specify how many image copies of the DBDS or area that DBRC is to maintain records of.

The GENMAX parameter specifies the maximum number of recovery generations (images) that DBRC maintains for the identified DBDS or DEDB area. Duplicate image copy data sets are not included in this number.

Use the RECOVPD parameter to maintain data for a certain period.

**Related reading:**

- See “Recovery period of image copy data sets and GENMAX” for more information about the RECOVPD parameter.
- See INIT.DBDS command in *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for more information on the REUSE, NOREUSE, GENMAX, and RECOVPD.

## Recovery period of image copy data sets and GENMAX

When you register a DBDS in the RECON data set, you specify the maximum number of image copy generations for DBRC to record with the GENMAX parameter of the INIT.DBDS command. When this number is exceeded, DBRC discards the information relating to the oldest image copy.

For example, if you take image copies daily and want to keep four days of back level image copies, specify a GENMAX of 4. However, in some emergency situations, you may take backup copies more frequently, possibly three in one day.

To prevent DBRC from discarding information relating to the earlier copies, specify the optional recovery period parameter (RECOVPD) to indicate the number of days you want information retained. If the GENMAX limit is reached, but the RECOVPD for the oldest image copy record has not expired, DBRC issues a warning message (DSP0065I), and does not discard the record. If the DSP0065I warning message appears frequently, you might need to tune the GENMAX or RECOVPD values with the CHANGE.DBDS command.

A *recovery period* is the minimum amount of time that information is maintained in the RECON data set. For example, if the recovery period of a DBDS or DEDB area is 14 days, DBRC maintains sufficient recovery-generation information for at least 14 days.

If both GENMAX and RECOVPD have been specified for a DBDS or DEDB area, DBRC considers both when deciding whether to reuse or delete an image copy data set.

The following table shows the results of GENJCL.IC processing when both GENMAX and RECOVPD have been specified for a DBDS or area defined with the REUSE parameter.

*Table 41. Results of GENJCL.IC processing when GENMAX and RECOVPD are specified with REUSE*

| Number of Image Data Sets | Number of In-Use Image Copies | Age of Oldest Image Copy | GENJCL Result               |
|---------------------------|-------------------------------|--------------------------|-----------------------------|
| =GENMAX                   | =GENMAX                       | <RECOVPD                 | Fail (DSP0063I)             |
| >GENMAX                   | =GENMAX                       | <RECOVPD                 | Avail IC DS used (DSP0065I) |
| =GENMAX                   | =GENMAX                       | >RECOVPD                 | Oldest IC DS reused         |
| >GENMAX                   | =GENMAX                       | >RECOVPD                 | AVAIL IC DS used            |



The following table shows the results of running an image copy utility when both GENMAX and RECOVPD have been specified for a DBDS or area defined with the NOREUSE parameter.

*Table 42. Results of GENJCL.IC processing when GENMAX and RECOVPD are specified with NOREUSE*

| Number of Image Copies | Age of Oldest Image Copy | Utility EOJ Result       |
|------------------------|--------------------------|--------------------------|
| =GENMAX                | >RECOVPD                 | Delete oldest image copy |
| =GENMAX                | <RECOVPD                 | No delete (DSP0064I)     |
| <GENMAX                | N/A                      | No delete                |

If you issue a CHANGE.DBDS command and specify new GENMAX and RECOVPD values that are less than the existing values, any used image copy data sets that are beyond the recovery period are deleted until the number of remaining image copy data sets equals the specified GENMAX value.

If you issue the DELETE.IC command, any specified image copy data set record is deleted regardless of RECOVPD or GENMAX.

**Related reference:**

 CHANGE.DBDS command (Commands)

## Reusing image copy data sets

DBRC enables you to reuse old image copy data sets. The REUSE parameter of the INIT.DBDS command, in addition to enabling you to define image copy data sets for future use, enables DBRC to reuse image copy data sets. To reuse the image copy data set means that the same name, volumes, and physical space are used for the new image copy.

**Restriction:** DFSMSdss fast replication, as invoked by the Database Image Copy 2 utility, does not support the REUSE parameter. Consequently, the information in this topic does not apply for fast replication image copies.

A run of one of the IMS image copy utilities automatically reuses the oldest image copy data set for a DBDS or area with the REUSE attribute when all of the following conditions are met:

- A number of image copy data sets equal to the current GENMAX value are recorded in the RECON data set. To see the current GENMAX value, use the LIST.DBDS command.
- The Database Image Copy utility or Online Database Image Copy utility used all image copy data sets for this DBDS. None of the image copy data sets is available.
- The oldest image copy is beyond the recovery period.
- A run of one of the IMS image copy utilities automatically uses an available image copy if one exists. If GENMAX has been reached and the oldest image copy is beyond the recovery period or there is no recovery period, the oldest image copy data set is made available.

When you use a GENJCL.IC command to generate the job for the Database Image Copy utility, the image copy data set that is to be reused is automatically selected. If the number of image copy data sets is less than the GENMAX value and all image copy data sets have been used, more image copy data sets must be defined

for the DBDS or area before running the Database Image Copy utility. The number of image copy data sets should be greater than the GENMAX value if you want to use a recovery period.

The Database Image Copy 2 utility can create up to four output image copy data sets. However, if you issue a GENJCL.IC command for a DBDS defined as REUSE, DBRC generates JCL for only one or two output copies (because you can define image copy data sets for only one or two copies). If you use the GENJCL.IC command for the Database Image Copy 2 utility to process a DBDS defined as REUSE and you want more than two output copies, modify the generated JCL before you run the job.

DBRC generates JCL so that the oldest DASD data set is always used for output for the image copy. DBRC has the same capability with tape volumes. However, you need to analyze your existing tape library techniques to make sure there is no conflict.

## **HSSP image copy**

The high speed sequential processing (HSSP) image copy (IC) process requires that a database be registered with DBRC and requires that an area be registered with the REUSE attribute for recycling the predefined IC data sets. Image copies that are created by HSSP are concurrent image copies.

## **HISAM copies (DFSURUL0 and DFSURRL0)**

Using the HISAM Reorganization Unload utility (DFSURUL0) to make backup copies of a database allows you to process an entire HISAM database in one pass (the image copy utilities deal with single data sets or areas). The unload utility (DFSURUL0) also reorganizes the database as it copies it.

Because the unload utility (DFSURUL0) reorganizes the database, before resuming normal online operations, the data set must be reloaded using the HISAM Reorganization Reload utility (DFSURRL0). The logging, which is done between the unload and reload, reflects the old data set organization.

When using the HISAM utility to make a backup copy, reload immediately, or the actual database and the backup database are mismatched.

The reload utility (DFSURRL0) notifies DBRC. The unload utility creates a reorganized copy of each data set. Then the reload utility reloads each data set from the reorganized copy, and through DBRC, creates a REORG and an IC record in the RECON data set for each data set. The IC record points to the data set that was output from the Unload utility and input to the Reload utility. After a database has been reorganized, a DBDS can be authorized only if an image copy of that data set has been created.

Updates of the database between unload and reload operations must be prevented. Updates of the database made after an unload operation but before a reload operation are wiped-out by the reload operation. In addition, the change records that are logged reflect the old organization, so that a subsequent recovery using those log records damages the database.

You can prevent access to a shared database during reorganization by using one of the following methods:

- From an online IMS subsystem, issue a global /DBRECOVERY command for the database that is to be reorganized. This prevents any subsequent authorizations except for reorganization and recovery. Ensure that recovery utilities do not run during the reorganization.
- Manually update the RECON data set by specifying the NOAUTH parameter of the CHANGE.DB command. This prevents any subsequent authorizations except for the reorganization and recovery utilities. Ensure that recovery utilities do not run during the reorganization. After the reorganization process is complete, manually update the RECON data set by specifying the AUTH parameter of the CHANGE.DB command for the database that was just reorganized.

## Nonstandard image copy data sets

You can create backup copies of DBDSs by using some means other than the Database Image Copy utility.

For example, you could make a copy of the volume on which a DBDS resides. DBRC does not record the existence of these nonstandard image copy data sets in the RECON data set automatically; use a NOTIFY.UIC command to inform the RECON data set of these data sets. If this information is not recorded in the RECON data set, DBRC might misinterpret subsequent information about changes to the DBDS. You cannot issue the NOTIFY.UIC command for a DBDS that is defined with the REUSE option.

Before recovering a DBDS or DEDB area with a nonstandard image copy data set, perform the following steps:

1. Close the database using /DBR (without NOFEOV). Load the data set from the nonstandard image copy (UIC) and record the event in the RECON data set (by issuing NOTIFY.RECOV with RCVTIME specified).
2. Apply the change records from the logs that were produced since the UIC (by running DBRC with USEDDBS or USEAREA for the GENJCL.RECOV command or DFSUMP DD DUMMY statement in the DBRC JCL).

DBRC does not allow DBRC to process any log that contains changes outside the recovery range because an image copy is not used for Step 2. The recovery range is defined by the time-stamp recovery record RECOV TO (image copy time) and RUNTIME values.

For more information about recovering a DBDS or DEDB area with a nonstandard image copy data set, see *IMS Version 11 Database Administration*.

## Frequency and retention of backup copies

Generally, the more frequently you copy, the less time recovery takes. The further back in time your old copies go, the further back in time you can recover. (Remember that program logic errors are sometimes not discovered for weeks.) Conversely, making each new copy requires work, and each old copy that you save uses additional resources.

Consider the following questions when making backup copies of databases:

- How frequently should I make new copies?
- How long should I keep old (back-level) copies?

There are no precise answers to these questions.

The only firm guidelines are these:

- If a database is composed of several data sets, be sure to copy all of the data sets at the same time.
- If you reorganize a database, immediately make a new backup copy of it. (This is not necessary for online DEDBs or HISAM reorganizations.)
- If you create a new database, immediately make a backup copy of it.
- If you perform a time-stamp recovery, make a backup copy for use in subsequent recoveries.

---

## Log record change accumulation

As IMS runs, the number of stored SLDSs or RLDSs grows. You can use these stored volumes to recover a lost or damaged database, but to use them in their raw form would be inefficient.

It would be inefficient because:

- Each SLDS or RLDS contains a record of activities of the entire system and all the data sets within all the databases. Yet when you are recovering a database, you are doing so for a single data set only. Thus, much of what is on the SLDS or RLDS does not apply.
- The SLDS or RLDS chronologically notes each change to any single record. If a record were changed 100 times since the last backup copy of the data set, the SLDS or RLDS would include 100 such notations. Yet, in recovery, you are only interested in the value the data had at the moment the data set was lost. The previous 99 changes are irrelevant.

The Database Recovery utility requires that change records be input in chronological order, but if a database is shared, the change records might be distributed among different log data sets in a way that makes their input to the utility impossible. A DBRC GENJCL.RECOV command or DB Recovery utility execution fails if this log data has not been properly merged. Such a failure is accompanied by a message informing you that a change accumulation must be run.

## Condensing the accumulated SLDS or RLDS (change accumulation)

The Database Change Accumulation utility offers you a way of sorting through your accumulated log data sets in advance, merging and condensing them.

This utility:

- Merges updates from different subsystems
- Picks out only those log records relating to recovery of databases
- Sorts these records by data set within a database
- Saves only the most recent state of each changed part of each data set.

The following figure illustrates how the Database Change Accumulation utility merges the relevant log data (a, b, and c) from 3 different SLDSs (A, B, and C), from one IMS system, into one change accumulation data set record.




Figure 53. What change accumulation does with redo data from divergent data streams

## When is change accumulation required?

Running the Database Change Accumulation utility is not required if, during the time period in question, only one system updated the database; using it periodically speeds any database recovery that becomes necessary. Alternately, you can run the Database Change Accumulation utility only when the need for recovery arises (just before running DB Recovery utility).

Change accumulation would not be required when non-concurrent data set update information exists in various logs. The database changes are received in the correct order if the logs are input serially to DB Recovery utility.

Change accumulation can be required when concurrent data set update information exists in various logs. The logs cannot be input to DB Recovery utility in a way that the change records are seen in the correct order.

**Related reading:** See *IMS Version 11 Database Utilities* for detailed instructions on running the Database Change Accumulation utility.

## Input to the database change accumulation utility

In addition to stored SLDSs and RLDSs, you can use a previous change accumulation data set and other IMS log volumes as input. The utility writes the accumulated changes in a new change accumulation data set.

You can specify all log volumes or a subset of log volumes as input to the Database Change Accumulation utility. When you specify a subset of log volumes,

DBRC determines whether the subset is complete for each DBDS or area. A subset of log volumes is complete for a DBDS or area when all of the following conditions are true:

- The first volume in the subset is the earliest volume, that could possibly have changes to the DBDS, that were not included in the last change accumulation or in the last image copy.
- The remaining volumes are in sequence.
- In a data sharing environment, logs from all updating subsystems containing changes and any open data streams for a DBDS are included.

The DBRC LIST.CAGRP command indicates whether the log subset for each DBDS of the change accumulation group is complete.

You can use the change accumulation data set as input to a later run of the Database Change Accumulation utility whether your subset of log volumes is complete or incomplete; however, you can use a change accumulation data set as input to DB Recovery utility only if it represents a complete log subset.

**Recommendation:** If DBRC is being used, perform the change accumulation process using the GENJCL.CA command. This command creates the correct JCL and includes all unprocessed log data sets. If you use your own JCL instead, verify the specifications for change accumulation before execution.

An image copy of the specified database data set is needed and must be identified to the RECON data set in order to create a valid starting point for change accumulation records.

All changes since the last valid image copy are collected by the utility. If a time-stamp recovery has occurred since the last image copy, the change accumulation that is created is invalid for use in future recoveries. No error messages are generated by GENJCL.CA or by the execution of the utility.

You can run the Change Accumulation utility with a valid log subset at any time to reduce data to a minimum.

When the most recent image copy is used as input to the Change Accumulation utility and that image copy is a concurrent image copy, changes already made to the database by active applications might be missing from the copy because the changes might not have been physically written to the data set. These changes, however, have been written to the log. In this case, it is necessary to go back to some earlier point in the logs to ensure that all changes are applied. How far to go back depends on the type of database and which image copy utility was used.

The point-in-time selected to start the Change Accumulation utility is called the *purge time*.

**Related reading:** See *IMS Version 11 Database Utilities* for more information on the Database Change Accumulation utility.

## Change accumulation groups

You can use the Database Change Accumulation utility to process some or all of your recoverable, registered databases as change accumulation (CA) groups. You can, for example, organize the DBDSs into the following groups:

- Application-associated databases

- Physical database clusters
- Logical database clusters
- Volatile (critical data) databases

You can add or delete members of a CA group after you have created it. A database can be a member in only one CA group. To move a member from one CA group to another CA group, you must first delete it from its current CA group and then add it to the new CA group.

**Recommendation:** Create an image copy of a CA-group member before you move that member to a new CA group. The change accumulation utility processes all logs for the member that are needed for recovery, based on the last good image copy.

A Change Accumulation group can contain from one to 32767 members, which must be registered DBDSs or areas.

Use the INIT.CAGRP command to create your Change Accumulation group. To modify your Change Accumulation group, use the CHANGE.CAGRP. A DBDS or area can belong to only one Change Accumulation group.

DBRC supports the Change Accumulation utility only by Change Accumulation group. Likewise, GENJCL.CA generates JCL only for a Change Accumulation group. The DBRC verification exit routine verifies for a whole Change Accumulation group.

You can run the Change Accumulation utility for DBDSs that do not belong to a Change Accumulation group even with DBRC=YES in effect. However, DBRC does not verify the input to the utility nor record its output.

## Defining change accumulation data sets for future use

You can use the REUSE keyword on the INIT.CAGRP command to inform DBRC that you want to define a “pool” of data sets to receive the output from the Change Accumulation utility. Define the data sets using the INIT.CA command. The number of change accumulation data sets that you can define, can equal up to the value of the INIT.CAGRP command GRPMAX parameter that defined the group.

If you define a Change Accumulation group with the REUSE parameter and also use a GENJCL.CA command to generate the job for the Database Change Accumulation utility, data set reuse can occur. When all available change accumulation data sets for this Change Accumulation group have been used and the maximum number of change accumulation data sets has been reached, the next run of the Database Change Accumulation utility for this group reuses the change accumulation data set containing the oldest change records. To reuse a change accumulation data set means that its data set name, volumes, and physical space are used as if they were for an empty change accumulation data set.

If you define the Change Accumulation group with NOREUSE:

- You must specify the output data set name for the utility either in the skeletal JCL member used for the GENJCL.CA command or in the JCL that you produce.
- When you run the Change Accumulation utility and the number of data sets specified by GRPMAX has been reached, DBRC deletes the record of the change accumulation data set (from the RECON data set) that contains the oldest



change records. The data set itself is not scratched. The data set must be manually scratched or monitored if you want to keep it, because DBRC no longer recognizes its existence.

**Related reading:** See *IMS Version 11 Database Utilities* and *IMS Version 11 Commands, Volume 1: IMS Commands A-M* for details on the Change Accumulation JCL specifications.

---

## DBDS group considerations

A DBDS group is a named collection of DBDSs or DEDB areas. DBRC can perform various operations by DBDS group so that you do not need to repeat the command for every member of the group.

You can specify DBDS groups on the following commands:

- GENJCL.IC
- GENJCL.OIC
- GENJCL.RECEIVE
- GENJCL.RECOV
- GENJCL.USER
- LIST.DBDS
- LIST.HISTORY

When you specify a DBDS group on a command, DBRC invokes that command for each member of the DBDS group. For example, you might have a DBDS group for a particular application, like payroll. When performing a time-stamp recovery, for example, all DBDSs of a particular application of a database must be recovered to the same point. If you specify a DBDS group on the GENJCL.RECOV command, you need only invoke the command once to recover all DBDSs.

You can also specify a CA group as a DBDS group. DBRC then executes the command for each member of the CA group.

You can define as many DBDS groups as you want. Up to 32767 DBDSs can be in a group. All DBDSs in a group must be registered in the RECON data set. A DBDS can belong to more than one DBDS group.

A database is an implied DBDS group for the GENJCL and LIST commands. It is unnecessary to define a DBDS group consisting of the DBDSs or areas of a single database. Specify the database name and omit the DD name to operate on the whole database.

The following commands affect the definition of a DBDS group:

- INIT.DBDSGRP
- CHANGE.DBDSGRP
- DELETE.DBDS
- DELETE.DBDSGRP
- LIST.DBDSGRP

DBDS groups can include ILDS (Indirect List Data Set) and index data sets.

**Related reading:** See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for the impact of these data sets on the GENJCL commands.



---

## DB groups

A DB group is a named collection of databases or DEDB areas. A DB group name can be specified in the /START, /STOP, /DBRECOVERY, and /RECOVER commands, instead of issuing these commands separately for each database or area.

Specifying a DB group name with these commands greatly reduces the number of times these commands must be issued. Use the DATAGROUP keyword to specify the DB group name.

DB groups can also include HALDB master and HALDB partition names. Be aware of the effects of a command issued using a DB group that has a HALDB master name and one or more of its partitions. See *IMS Version 11 Database Administration* for more information.

You can define as many DB groups as you want. Up to 32767 databases or areas can be in a group. A database or area can belong to more than one DB group and need not be registered in the RECON data set.

**Recommendation:** Use a database group whenever possible, even though a DBDS group can be used as a DB group. Processing a DBDS group as a DB group entails increased overhead.

DB groups are a form of a DBDS group, so they are stored in the RECON data set using the DBDS group record. The following commands affect the definition of a DB group:

- INIT.DBDSGRP
- CHANGE.DBDSGRP
- DELETE.DBDSGRP
- LIST.DBDSGRP

---

## DBRC groups

When more than one DBRC accesses the same RECON data set in an IMSplex environment, they are considered a DBRC group.

The group:

- Is identified by a DBRC group ID.
- Can access the shared RECON data set in serial or parallel mode.
- Is supported in IMS Version 10 and later.

A DBRC group ID can be specified in the CHANGE.RECON command. You can view the DBRCGRP in the QUERY IMSPLEX command in the subtype column.

**Restriction:** DBRCs in the same IMSplex that use different RECON data sets must use different DBRC group IDs when parallel RECON data set access is being used.

**Recommendation:** If serial RECON access is being used in the IMSplex, DBRCs in the same IMSplex that use different RECON data sets can use the same DBRC group ID, but doing so is not a good idea. When one RECON data set encounters an error, automatic RECON loss notification processing is performed by all DBRC instances in the IMSplex with the same DBRC group ID, thus adding processing

overhead to the DBRC instances that are not using the RECON data set in error. Using different DBRC group IDs avoids this additional processing overhead.

---

## DBRC considerations for HALDB Online Reorganization

The following topics describe the DBRC considerations with respect to HALDB online reorganization.

The following DBRC commands support HALDB online reorganization:

- CHANGE.DB
- GENJCL.IC and GENJCL.OIC

You can generate the JCL for the GENJCL.IC and GENJCL.OIC commands in advance and not execute the JCL immediately. Therefore, the active set of DBDSs might change in the interim time period.

GENJCL no longer generates DD statements for the HALDB DBDS that is being copied because dynamic allocation is recommended for the Database Image Copy utility and the Database Image Copy 2 utility. The SYSIN control statement identifies which DBDS to copy: either the A-through-J or the M-through-V set of data sets. The image copy utility then copies the corresponding active set of data sets: A-through-J or M-through-V.

- INIT.DB
- NOTIFY.IC, NOTIFY.REORG, and NOTIFY.UIC

You cannot take an image copy if a HALDB online reorganization cursor is active (OLREORG CURSOR ACTIVE =YES) or has been terminated. The NOTIFY.IC and NOTIFY.UIC commands fail if they attempt to record an image copy while the HALDB online reorganization cursor is active.

HALDB online reorganization relies on a number of records to track information and activities related to online reorganization. This section includes descriptions of some of the records impacted by HALDB online reorganization:

- HALDB master record
- Partition DB record

The DB record for a HALDB partition contains information for HALDB online reorganization. For a complete description of the fields in the DB record for a HALDB partition, see “Fields present in a listing of a RECON by record type” in *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*. The DB record includes:

- A field that indicates whether the database can be reorganized online (ONLINE REORG CAPABLE=). This field is also in the HALDB master DB record. This field cannot be reset if:
  - Online reorganization is active
  - Any of the M-through-V and Y data sets are active
  - Any partition for the HALDB is authorized

If these conditions do not exist and the field can be reset, all of the M-through-V and Y data set information in the RECON data set is deleted.

- A field that indicates whether the HALDB online reorganization cursor is active (OLREORG CURSOR ACTIVE=).
- The SSID of the IMS that is currently performing a HALDB online reorganization (OLRIMSID=).
- A field that indicates which set or sets of DBDSs is or are active for the partition (ACTIVE DBDS=).

- Counters showing the number of cursor-active status changes (OLR ACTIVE HARD COUNT=) and cursor-inactive status changes (OLR INACTIVE HARD COUNT=) that are not yet hardened at the RSR tracking site.

- DBDS record

The M-through-V and Y DBDS records are defined in the RECON data set with the same attributes as the corresponding A-through-J, L, and X DBDS records. The M-through-V and Y DBDS records are added to the same CAGROUP as the A-through-J and X DBDS. Similarly, the M-through-V and Y DBDS records are added to each DBDS group where the A-through-J DBDS records are found.

The exceptions to the similarities in attributes include DD names, data set names, and the DSID. The DSID of the M-through-V and Y data set is the same as the corresponding A-through-J, L, and X data set, except that its X'80' bit is set on. For the RECON data set list purposes, the X'80' bit is ignored.

The following table shows the data set ID (DSID) DCB number associated with the A-through-J, L, and X data sets and the M-through-V and Y data sets.

*Table 43. DSID DCB numbers for the A-J, L, and X and the M-V and Y data sets*

| DSID DCB<br>Number | A-J, L, and X Data Sets |       | M-V and Y Data Sets |       |
|--------------------|-------------------------|-------|---------------------|-------|
|                    | PHIDAM                  | PHDAM | PHIDAM              | PHDAM |
| 1                  | A                       | A     | M                   | M     |
| 2                  |                         |       |                     |       |
| 3                  | L                       | L     |                     |       |
| 4                  |                         | B     |                     | N     |
| 5                  | X                       | C     | Y                   | O     |
| 6                  | B                       | D     | N                   | P     |
| 7                  | C                       | E     | O                   | Q     |
| 8                  | D                       | F     | P                   | R     |
| 9                  | E                       | G     | Q                   | S     |
| 10                 | F                       | H     | R                   | T     |
| 11                 | G                       | I     | S                   | U     |
| 12                 | H                       | J     | T                   | V     |
| 13                 | I                       |       | U                   |       |
| 14                 | J                       |       | V                   |       |

- REORG record

The REORG record for a HALDB partition DBDS contains a stop time for HALDB online reorganization. To run an offline reorganization of a database, you must take it offline by issuing a /DBR command. Therefore, only the start time of the online reorganization is significant. With HALDB online reorganization, the database is online, so the duration of the process is significant. The stop time of the online reorganization is recorded in the REORG record (STOP=*timestamp*), with information that indicates that it is an online reorganization and indicates whether the online reorganization can be used as input to recovery.

**Related reading:** For more information on Fields present in a REORG record, see *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*.

A REORG record is recorded under both the A-through-J, L, and X DBDS and the M-through-V and Y DBDS. The REORG record for the DBDS that was the output data set indicates that it can be used as input for recovery. The REORG

record for the DBDS that was the input data set indicates that it cannot be used for recovery after online reorganization completes.

## RSR tracking system

Ownership of the HALDB online reorganization is not recorded in the tracking RECON data set.

The RECON data set at the tracking site always reflects the current status of the covered databases. Counters keep track of outstanding HALDB online reorganizations that have been recorded in the RECON data set at the tracking site but not yet applied to the databases.

## Maintaining recovery-related records affected by online reorganization

For recovery purposes, the A-through-J and M-through-V data sets are treated in the RECON data set as one logical set of DBDSs.

The recovery information in the two DBDS headers is maintained as if they are one logical DBDS. For example, the values for GENMAX, RECOVPD, and CAGRP are the same. Any changes that are made to the active DBDS are also made to the inactive DBDS.

There is a one-to-one relationship between each partition data set; for example, A correlates to M, B correlates to N, C correlates to O, and so on. Also, the M-through-V data sets are defined in the RECON data set with the same attributes as their corresponding A-through-J data sets. You can perform online reorganization on any of these data sets (A-through-J, L and X or M-through-V and Y) at any time, depending on which set is active before the online reorganization starts. Therefore, data sets A and M are one logical data set during an online reorganization.

After an image copy is taken, DBRC removes any extraneous recovery-related records (IMAGE, ALLOC, RECOV, and REORG) from the RECON data set. The extraneous records are determined by the maximum number of recovery generations (GENMAX) settings and the optional recovery period (RECOVPD) settings for the DBDS.

For example, assume that the GENMAX value for data set A (and M) is 2, and that data set A has two image copies. When data set M becomes the active data set and an image copy is taken, the image copy cleanup process deletes the oldest image copy from data set A, including any of its extraneous recovery records. Another image copy of data set M causes the rest of data set A's image copy and recovery records to be deleted. Note that data set A becomes unrecoverable if its image copy and recovery records are deleted.

## Online reorganization process

When online reorganization is started, the ID of the IMS that is responsible for the reorganization is assigned to the partition.

After the cursor-active status is recorded in the RECON data set and copying begins between the A-through-J and M-through-V data sets, these changes also occur:

- The DB record for the partition indicates that online reorganization is active (OLREORG CURSOR ACTIVE =YES).
- The status of the active DBDSs shows that both data sets are active (ACTIVE DBDS=A-J and M-V).
- REORG records that show the start time of the online reorganization are written. The stop time is zero.
- The REORG record associated with the output data set now indicates that this data set can be used as input for recovery.
- The REORG record associated with the input data set does not indicate that this data set can be used for recovery. DBRC uses this record to ensure that this data set cannot be recovered after online reorganization completes.
- The M-through-V and Y data sets are defined in the RECON data set if they do not already exist.
- The M-through-V and Y DBDSs are added to the same change accumulation group as the A-through-J, L, and X DBDSs, if a change accumulation group is assigned.
- The M-through-V and Y DBDSs are added to any DBDS group where the A-through-J, L, and X DBDSs are found.

After these changes take place, the online reorganization must complete; otherwise, you must run an offline reorganization to reset the HALDB online reorganization settings.

You cannot start online reorganization if it is already running for a partition. However, another IMS can take over the online reorganization if it is terminated; that is, when there is no IMS ID assigned.

When the online reorganization process completes successfully, these changes occur:

- The REORG records are updated to include the HALDB OLR stop time.
- The DB record for the partition is changed to indicate that online reorganization is not active (OLREORG CURSOR ACTIVE = NO).
- The status of the active DBDSs is switched to the newly reorganized data sets.
- The online reorganization IMS ID is cleared.

**Recommendation:** Create an image copy of the HALDB partition DBDSs as soon as possible after a successful online reorganization.

**Restriction:** An image copy of a DBDS cannot be taken while the online reorganization cursor is active (OLREORG CURSOR ACTIVE =YES). DBRC causes the image copy to fail in this situation.

---

## Chapter 41. Initializing and maintaining the RECON data sets

DBRC records recovery-related information in a pair of key-sequenced data sets (KSDSs) called the REcovery CONTROL (RECON) data set. DBRC uses two RECON data sets to increase availability and recoverability. They contain identical information. The data sets are identified by the DD names RECON1 and RECON2.

If you define only two RECON data sets and an error occurs on one of them during operation, the current jobs continue using the remaining one. New jobs cannot start unless your RECON setting allows a job to start with only one active RECON data set. If you want to continue operations in dual mode, you can define a third RECON data set (RECON3). DBRC does not use this spare data set unless an error occurs on one of the two active RECON data sets. Then, DBRC copies the good RECON to the spare data set (RECON3), which then becomes active (thus maintaining RECON dual-mode operation).

The following figure shows the recommended three RECON data set operating configuration.




Figure 54. DBRC dual RECON set up with a spare

The RECON data sets are critical resources for both DBRC and IMS. If both RECON data sets are lost, DBRC abnormally terminates rather than compromising database integrity. IMS cannot continue processing transactions without a viable RECON data set, so IMS also abnormally terminates.

---

### Planning considerations for the RECON data sets

Define your RECON data sets specifying VSAM SHAREOPTION(3,3). Ideally, each data set should be on a different device, channel, control unit, and catalog. The RECON data sets should also be of different sizes.

Each RECON data set is a VSAM KSDS with a 32-byte key.

#### Recommendations:

- When defining the RECON data sets using Access Method Services (AMS):

- Use the same index control interval (CI) size and data CI size for all the RECON data sets.
- Ensure that the data CI size specified exceeds the specified index CI size by at least 2048 bytes. Failure to do so can seriously degrade your DBRC performance.
- Use secondary allocation to safeguard against a RECON data set becoming full.

**Related concepts:**

“Avoiding RECON data set space problems” on page 542

## Initializing the RECON data sets for the first time

After allocating the RECON data sets, use the INIT.RECON command to initialize the RECON data set.

This command is only valid for an empty, un-initialized RECON data set. If the initialization job fails, delete and redefine the data sets before rerunning the job.

## Avoiding RECON data set contention problems

Maximum RECON data set availability depends on eliminating deadlock situations and minimizing situations in which more than one RECON data set becomes simultaneously unavailable.

For maximum availability, each RECON data set must:

- Have different space allocations. The spare data set must be at least as large as the largest RECON data set.
- Be on different devices.
- Be on different channels.
- Be in different user catalogs.

## Accessing the RECON data sets in parallel mode

You can access the RECON data sets in parallel mode to avoid data set deadlock situations.

The following terms are used in the topics related to parallel mode:

### DBRC group

One or more DBRC instances that share a single RECON data set in an IMSplex environment. DBRC groups are supported only in IMS Version 10 and later.

### DBRC group ID

The identifier of a DBRC group

### DBRC request

Each action asked of DBRC (command process, query, and update) by IMS, DBRC application programs, and jobs.

DBRC uses the Transactional VSAM function of DFSMS to provide parallel access to the RECON data set. Transactional VSAM uses:

- The two-phase commit and backout service provided by z/OS Resource Recovery Services to ensure that either all of the changes made by a transaction are hardened or all of the changes made by a transaction are backed out.
- A coupling facility (CF) for buffer caching, locking, and logging services.

Use the following commands to set the RECON data set access mode to either serial or parallel:



- For a new RECON data set, use a `INIT.RECON ACCESS(SERIAL or PARALLEL)` command.
- For an existing RECON data set, use a `CHANGE.RECON ACCESS(SERIAL or PARALLEL)` command.

Before parallel RECON data set access can be enabled, both of the following conditions must be true:

- All DBRC instances that are accessing the RECON data set belong to a DBRC group with a unique DBRC group ID.
- All DBRC instances that are accessing the RECON data set are registered to Structured Call Interface (SCI), similar to the situation where RECON Loss Notification is involved. DBRC registers with SCI using the group ID as the member subtype.

### Planning for parallel RECON data set access:

Set up the Structured Call Interface and Transactional VSAM to prepare for parallel RECON data set access.

You must perform the following tasks to prepare for parallel RECON access:

- Set up the Structured Call Interface (SCI).
  - Define a Common Service Layer (CSL) and configure the IMSplex environment such that there is a Structured Call Interface on each z/OS image from which IMS systems and batch jobs will access the RECON data set.
  - Tailor the DBRC SCI Registration exit routine, `DSPSCIX0`, as necessary to specify the correct IMSplex name for the RECON data set. Ensure that your exit routine is coded to work with the IMS Version 10 and later `DSPSCIX0` extended interface.
- Set up Transactional VSAM:
  - Update the `IGDSMSxx` member in `SYS1.PARMLIB` to support Transactional VSAM.

**Recommendation:** Specify the `DEADLOCK_DETECTION` in `IGDSMSxx` so that detection times for both local and global detection are less than 2 seconds.

- Set up the SMS configuration to enable Transactional VSAM processing.
- If the RECON data sets are not SMS-managed, convert them to be SMS-managed.
- Specify a storage class for each RECON data set to assign the data set to a coupling facility (CF) cache structure.
- Define at least one CF cache structure for the SMS configuration.
- Define to z/OS the CF lock structure, `IGWLOCK00`.
- Define a primary system log stream for each instance of Transactional VSAM.
- Define a secondary system log stream for each instance of Transactional VSAM.

**Related reading:** For more information about configuring SMS and setting up a CF structure for Transactional VSAM, see the *z/OS DFSMS Storage Administration Reference (for DFSMSdfp, DFSMSdss, and DFSMShsm)*. For more information about `SYS1.PARMLIB` changes and log stream definition for Transactional VSAM, see the *z/OS TSO/E Customization*.



- If you use a RECON I/O exit routine (DSPCEXT0), ensure that the exit routine works with the IMS Version 10 and later extended RECON I/O exit interface.
- Ensure that users that issue the CHANGE.RECON.ACCESS command, using either the batch command or through the DBRC API, have the following authority permissions:
  - Data set ALTER authority to maintain the LOG parameter on the DEFINE CLUSTER command
  - UPDATE authority to the specified data set and to the RACF class profile STGADMIN.IGWSHCDS.REPAIR
- Ensure that read-only jobs have READ authority to the RACF class profile STGADMIN.IGWSHCDS.REPAIR.
- Except for jobs that will run in online DBRC regions, any jobs that are not read-only must have UPDATE authority to the class profile STGADMIN.IGWSHCDS.REPAIR. If UPDATE authority is not allowed, these regions might not be able to recover from errors that involve shunted I/O.
- Ensure that all DBRC instances that share the RECON data set use the same IMSplex name for the Parallel RECON Access and Automatic RECON Loss Notification functions.

If you activate the Automatic RECON Loss Notification and Parallel RECON Access functions within the same IMSplex, you must use the CHANGE.RECON IMSPLEX() command to ensure that all DBRCs in the IMSplex are using the same IMSplex name as is specified in the RECON data set. If you use the DBRC SCI Registration exit routine (DSPSCIX0) or use the IMSPLEX EXEC parameter before issuing the CHANGE.RECON IMSPLEX() command, message DSP1136A is issued and subsequent jobs fail due to an unavailable RECON data set.

### **Avoiding RECON data set deadlock situations for serial access mode**

To eliminate deadlocks (if you are using implementing hardware reserves for the RECON data sets), the RECON data sets must be the only objects cataloged in their respective catalogs and be on the same device as their catalogs. When the RECON data sets are accessed, an enqueue on the RECON data sets can result, followed by an enqueue on the catalog. When the RECON data sets and catalog are on the same device, the possibility of conflicts with another job enqueueing the devices in reverse order is eliminated.

Give special consideration to the placement of RECON data sets that are shared among multiple processors. During a physical open, DBRC reserves RECON1, RECON2, and RECON3. DBRC determines which are available and which are Copy1, Copy2, and spare. DBRC then closes and dequeues the spare (if it exists) and any unusable RECON data sets. So, during the use of DBRC, two RECON data sets are reserved most of the time. DBRC always reserves both RECON data sets in this order: RECON1 and RECON2. If RECON1 and RECON2 are specified consistently throughout jobs, DBRC does not encounter deadlock.

However, other jobs that reserve multiple volumes can cause deadlock if any of the volumes also contain a RECON data set.

**Recommendation:** To eliminate contention and facilitate recoverability, the RECON data sets should be the only type of data set on their respective devices if using implementing hardware reserves for the RECON data sets.

Access to the RECON data sets must be controlled to avoid deadlock. The following z/OS macros are used to control access to the RECON data sets:

RESERVE, GRS, OBTAIN, and DEQ. The z/OS macros, DFP Record Management Services, and RECON data set serialization strategies are discussed in the following list.

- RESERVE

DBRC issues the z/OS RESERVE macro to serialize access to each RECON data set. DBRC keeps the RECON data sets reserved until it completes its processing. The more RECON records DBRC must examine or change, the longer it holds the RECON data sets. The RESERVE macro serializes access to a resource (a data set on a shared DASD volume) by obtaining control of the volume on which the resource resides to prevent jobs on other systems from using any data set on the entire volume. This reserve is done under the major name DSPURI01 and has a scope of SYSTEMS.

Batch jobs will serialize on another resource name first, before issuing a RESERVE for the RECON data sets. The resource name for batch is DSPURI02 and has a scope of SYSTEM. Each job gets control of the resource, based on the position of the task's request and whether the request was exclusive or share control. The queue is not ordered by the priority of tasks. The effect of this serialization of batch jobs is that an IMS online region never has to wait for more than one batch job to complete before gaining access to the RECON data set.

- Global Resource Serialization (GRS)

The GRS macro provides a method of converting a RESERVE request into an ENQ request. The ENQ, DEQ, and RESERVE macros identify a resource by its symbolic name. The symbolic name has three parts:

- Major name (qname)
- Minor name (rname)
- Scope (which can be STEP, SYSTEM, or SYSTEMS)

For example, on an ENQ or DEQ macro, a resource might have a symbolic name of APPL01,MASTER,SYSTEM. The major name (qname) is APPL01, the minor name (rname) is MASTER, and the scope is SYSTEM.

When an application uses the ENQ, DEQ, and RESERVE macros to serialize resources, global resource serialization uses resource name lists (RNLs) and the scope on the ENQ, DEQ, or RESERVE macro to determine whether a resource is a local resource or a global resource. Global resource serialization identifies each resource by its entire symbolic name. For example, a resource that is specified as A,B,SYSTEMS is considered a different resource from A,B,SYSTEM or A,B,STEP because the scope of each resource is different. To ensure that resources are treated as you want them to be without changes to your applications, global resource serialization provides three resource name lists (RNLs):

- SYSTEMS EXCLUSION RNL

The SYSTEMS exclusion RNL contains a list of resources that are requested with a scope of SYSTEMS that you want global resource serialization to treat as local resources.

- RESERVE CONVERSION RNL

The RESERVE conversion RNL contains a list of resources that are requested on RESERVE macros for which you want global resource serialization to suppress the RESERVE.

- SYSTEM INCLUSION RNL

The SYSTEM inclusion RNL contains a list of resources that are requested with a scope of SYSTEM that you want global resource serialization to treat as global resources.

**Related reading:** See the following publications for more information about GRS and the z/OS RESERVE, DEQ, and ENQ macros:

- *z/OS MVS Programming: Assembler Services Reference Vol 1*
- *z/OS MVS Planning: Global Resource Serialization*

- **OBTAIN**

DBRC uses a VSAM DADSM (Direct Access Device Storage Management) OBTAIN request to the FORMAT-4 DSCB (the VTOC) to force an I/O that insures that DBRC really has the RECON data set reserved in a multisystem environment.

The OBTAIN macro is not issued if the Synchronous RESERVE (SYNCHRES) option is in effect. The Synchronous RESERVE feature allows specification upon installation of whether the hardware RESERVE should be obtained for a device prior to granting a global resource serialization ENQ.

The SYNCHRES option can be activated through either the GRSCNFxx parmlib member or the SETGRS operator command. The GRSDEF statement of GRSCNFxx contains the SYNCHRES (Y | N) parameter. The default value for SYNCHRES is NO. During normal system operation, the operator can modify the setting of SYNCHRES by issuing the SETGRS command. Activate SYNCHRES by issuing SETGRS SYNCHRES=YES and deactivate it by issuing SETGRS SYNCHRES=NO.

- **DEQ**

DBRC releases the RECON data sets by using the z/OS DEQ macro.

- **DFP Record Management Services**

DBRC uses VSAM services to retrieve, manipulate, and store the RECON data set records. These records have a 32-byte record key.

- **RECON data set serialization strategies**

**Recommendations:**

- In a GRS Star configuration, a RESERVE CONVERSION RNL should be implemented for DSPURI01 if all systems accessing the RECON data sets are within the sysplex (or GRSPlex).

**Note:** If the RECON data sets are accessed by systems that are outside of the sysplex, the reserve must not be converted. A SYSTEMS EXCLUSION RNL must be implemented instead.

If you implement GRS RNL CONVERSION by adding the QNAME for the RECON data set, DSPURI01, to the conversion list, the hardware reserve is eliminated and replaced by a GRS enqueue that is communicated to all other sharing z/OS systems.

Other data sets on the same DASD volume can be used while the RECON data set is reserved; this is the benefit of performing the RNL conversion. GRS RNL conversion uses CPU and storage and affects system performance positively. The performance (in terms of least CPU time used, least storage used, and least elapsed time) is best using this option in a GRS STAR configuration.

To implement this method, follow these steps:

1. Add the RECON data set QNAME to the RESERVE CONVERSION RNL. For example:  

```
RNLDEF RNL(CON) TYPE(GENERIC) QNAME(DSPURI01)
```
2. Consider carefully the placement of the following VSAM QNAMEs: SYSZVVDS and SYSIGGV2.

**Related reading:** For more information on implementing GRS RNL CONVERSION, see *z/OS MVS Planning: Global Resource Serialization*.

- In a GRS Ring configuration, a SYSTEMS EXCLUSION RNL should be implemented for DSPURI01 so that the RECON is serialized by hardware reserve.

If you implement GRS SYSTEMS EXCLUSION RNL, then GRS does not perform global serialization and the RESERVE macro is issued. This might work well provided that the RECON data set is located on a DASD volume that does not contain other data sets that are needed by other z/OS systems.

To implement this method, follow these steps:

1. Add the RECON data set QNAME to the SYSTEMS EXCLUSION RNL.

For example:

```
RNLDEF RNL(EXCL) TYPE(GENERIC) QNAME(DSPURI01)
```

2. Carefully consider the placement of the following VSAM QNAMEs; SYSZVVDS and SYSIGGV2.

**Related reading:** The *z/OS MVS Planning: Global Resource Serialization* provides more information about VSAM QNAMEs.

## Allocating the RECON data sets

For both online and batch DBRC jobs, you can allocate the RECON1, RECON2, and RECON3 data sets with JCL, or you can let DBRC dynamically allocate them.

When dynamically allocating the RECON data set, omit the DD statements for RECON1, RECON2, and RECON3.

Use the TYPE=RECON macro statement of the IMS DFSMDA dynamic allocation macro to establish three dynamic allocation parameter lists in IMS.SDFSRESL. When multiple processors access the same RECON data sets, keep the IMS.SDFSRESL information pertaining to dynamic allocation parameters in synchronization on all processors.

**Related reading:** See *IMS Version 11 System Definition* for information about the DFSMDA macro.

DBRC always allocates RECON data sets with DISP=SHR.

Although JCL allocation and dynamic allocation are both valid methods for allocating RECON data sets, JCL allocation should be used only in a controlled test.

**Recommendation:** Use dynamic allocation for your production system and all other test or development environments.

The principal advantages of dynamically allocating the RECON data set are:

- All DBRC jobs automatically use the correct and current RECON data sets, and no JCL statements are left to become outdated.
- You can reorganize and restore RECON data sets, in case of error, without having to shut down online IMS systems.

**Requirement:** When accessing the RECON data sets in parallel mode, they must be managed by DFSMS.

## Avoiding RECON data set space problems

Allocate the RECON data sets with different amounts of space so that if one becomes full, the system can continue using the other RECON data set while you provide a replacement. Allocate secondary extents for the RECON data set when you define space for it.

If one RECON data set becomes full during online operation, IMS deallocates it. DBRC responds by copying and reorganizing the good RECON data set to a spare RECON data set, if one is available. If no spare RECON data set is available, the system runs in single-RECON mode.

When all active subsystems have deallocated the failing RECON data set, you can delete and redefine it offline using AMS. See “Replacing a discarded RECON data set” on page 561 for information on replacing a discarded RECON data set. If you are in single mode, and a spare RECON data set is available, the next time DBRC accesses the RECON data set, it automatically enters dual-RECON mode. You do not have to enter the CHANGE.RECON command with the DUAL or REPLACE option.

## Creating a RECON data set

Create a RECON data set using the DEFINE CLUSTER command.

The keywords that are recommended when defining a RECON VSAM KSDS are discussed in the following list. Information on all the keywords can be found in the *z/OS DFSMS Access Method Services for Catalogs*.

### CONTROLINTERVALSIZE

The values used with this keyword affect the total amount of storage used by DBRC for VSAM and internal buffers. DBRC uses the Local Shared Resources (LSR) option of VSAM to process the RECON data sets. If the number of index and data buffers created by DBRC is allowed to default, the amount of storage used for RECON buffers is:

$$(60 \times \text{index\_ci\_size}) + (120 \times \text{data\_ci\_size})$$

This amount of storage is used when the index and data CI sizes are the same for all RECON data sets. You can change the default number of index or data buffers used by DBRC in an online or batch environment with the DSPBUFFS Buffer Size Specification Facility.

DBRC uses one set of internal buffers for each RECON data set to build logical records from smaller segments. The size of these buffers depends initially upon the CI size or the VSAM maximum record size (LRECL), whichever is smaller. After the initial allocation of these buffers, the logical record buffers grow over time as the largest logical record on the RECON data set grows. Because only the initial allocation of buffers is determined by the VSAM data set definition, the storage allocated to internal buffers cannot ultimately be affected by changes to either the CI or the record size.

DBRC divides its own records into segments, each of which is always smaller than a single control interval and which is seen by VSAM as a complete physical record. VSAM record spanning is not used. Segmenting allows a logical RECON record to be as large as 16 MB independent of the VSAM RECORDSIZE parameter.

**Recommendation:** Initially set your minimum CI size to a minimum of eight KB. The allowable CI size is affected by the value you select for RECORDSIZE.

Also, ensure that the smallest data CI size exceeds the largest index CI size by at least 2048 bytes. Failure to do so can seriously degrade your DBRC performance.

**Related reading:** See *IMS Version 11 Exit Routines* for further details about using the Buffer Size Specification Facility (DSPBUFFS).

#### **CYLINDERS**

Specifies the amount of space to allocate to the cluster.

#### **FREESPACE**

The default values of FREESPACE(0 0) must not be used. While you are entering initial information in the RECON data set, you must specify a high control-interval percentage (for example, 70%) as free space. Later, you can lower the percentage with an Access Method Services (AMS) ALTER command.

#### **INDEXED**

Specifies that the cluster is defined for key-sequenced data.

#### **KEYS**

KEYS(32 0) is required.

#### **LOG**

When parallel RECON access is used, let DBRC manage the LOG parameter on the AMS DEFINE CLUSTER command. When parallel RECON access is enabled, DBRC sets the value of the LOG parameter to UNDO (meaning that the data sets are recoverable) and sets the ACCESS= field in the RECON header to PARALLEL. When parallel RECON access is disabled, DBRC sets the value of the LOG parameter to NONE and sets the ACCESS= field in the RECON header to SERIAL. You should specify the appropriate value for the LOG value whenever you delete and redefine a RECON data set. For example, if you are using parallel access, you should specify LOG(UNDO)

#### **NAME**

Defines the cluster's entry name and is required.

#### **NONSPANNED**

See RECORDSIZE.

#### **RECORDSIZE**

DBRC always writes physical VSAM records that are less than the CI size in length, even though the logical RECON records can be as long as 16 MB.

**Recommendation:** Set the maximum record size to be, at most, equal to the CI size minus seven bytes and ensure that NONSPANNED is specified. For example, if the CI size was defined as 8192 bytes, then RECORDSIZE (4086,8185) is appropriate.

#### **SHAREOPTIONS**

SHAREOPTIONS(3 3) must be specified. The first value is required with single-host processors. Both values are required with multiple-host processors.

#### **SPEED**

This is recommended because the initial load is faster.

#### **NOWRITECHECK**

Avoid using WRITECHECK. It can degrade the RECON data set I/O performance. Using dual RECON data sets eliminates the need for WRITECHECK.



**Recommendation:** Do not use authorization keywords because frequent operator prompting results.

### Security considerations for the RECON data sets

Jobs that update the RECON data sets must have UPDATE authority to the RECON data sets.

For jobs that require read-only access, READ authority should be used. With parallel access, the job that issues the CHANGE.RECON ACCESS command must have ALTER authority. When parallel access is being used, the online-IMS DBRC regions should have ALTER authority.

For more information about preserving the integrity of the RECON data sets, see “DBRC security” on page 563.

---

## Initially accessing the RECON data sets

When a job needs to read the RECON data set, the RECON data set must be opened. When three RECON data sets exist, DBRC determines which two are the active RECON data sets and which is the spare data set.

The following table shows how this determination is made. The columns in the table have the following meanings:

- In the DD statement column, RECONA, RECONB, and RECONC represent the RECON1, RECON2, and RECON3 DD statements in no particular sequence.
- The data set status column indicates the status of the data set during open time.
- The data set use column indicates how DBRC assigns the data set.

*Table 44. Determining which RECON data sets are accessed.*

| Case | DD statement | Data set status | Data set use | DBRC selection criteria                            |
|------|--------------|-----------------|--------------|----------------------------------------------------|
| 1    | RECONA       | Create Mode     | Copy1        | Not selected until INIT.RECON command is specified |
|      | RECONB       | Create Mode     | Copy2        | Not selected until INIT.RECON command is specified |
|      | RECONC       | Create Mode     | Spare        |                                                    |
| 2    | RECONA       | RECON           | Copy1        | Current copy of RECON                              |
|      | RECONB       | Create Mode     | Copy2        | Produced from Copy1                                |
|      | RECONC       | Create Mode     | Spare        |                                                    |
| 3    | RECONA       | RECON           | Copy1        | Current copy of RECON                              |
|      | RECONB       | RECON           | Copy2        | Current copy of RECON                              |
|      | RECONC       | Create Mode     | Spare        |                                                    |
| 4    | RECONA       | RECON           | Copy1        | Current copy of RECON                              |
|      | RECONB       | RECON           | Unused       | Older copy than RECONA                             |
|      | RECONC       | Create Mode     | Copy2        | Produced from Copy1                                |
| 5    | RECONA       | RECON           | Copy1        | Current copy of RECON                              |
|      | RECONB       | RECON           | Unused       | Older copy than RECONA                             |
|      | RECONC       | RECON           | Unused       | Older copy than RECONA                             |

Table 44. Determining which RECON data sets are accessed (continued).

| Case | DD statement | Data set status | Data set use | DBRC selection criteria                            |
|------|--------------|-----------------|--------------|----------------------------------------------------|
| 6    | RECONA       | RECON           | Copy1        | Current copy of RECON                              |
|      | RECONB       | RECON           | Copy2        | Current copy of RECON                              |
|      | RECONC       | RECON           | Unused       | Older copy than RECONA                             |
| 7    | RECONA       | RECON           | Copy1        | Current copy of RECON                              |
|      | RECONB       | Create Mode     | Copy2        | Produced from Copy1                                |
| 8    | RECONA       | Create Mode     | Copy1        | Not selected until INIT.RECON command is specified |
|      | RECONB       | Create Mode     | Copy2        | Not selected until INIT.RECON command is specified |
| 9    | RECONA       | RECON           | Copy1        | Current copy of RECON                              |
|      | RECONB       | RECON           | Copy2        | Current copy of RECON                              |
| 10   | RECONA       | RECON           | Copy1        | Current copy of RECON                              |
|      | RECONB       | RECON           | Unused       | Older copy than RECONA                             |
| 11   | RECONA       | Create Mode     | None         | Discontinue processing                             |
| 12   | RECONA       | RECON           | Copy1        | Current copy of RECON                              |

Case 4 is the situation where two RECON data sets are available, but one is now out of date. DBRC does not use the out-of-date RECON data set. Instead, it copies the up-to-date RECON data set to the spare data set.

Only one RECON data set is available in cases 5, 10, and 12. If you have specified the STARTNEW parameter of the INIT.RECON or CHANGE.RECON command, processing continues with one RECON data set. Otherwise, processing ends. If another DBRC instance is active and parallel RECON access is enabled, the new DBRC must be able to access the same data sets (both in name and active count) as the existing DBRC regardless of the STARTNEW parameter. If the new DBRC does not access the same data sets, a WTOR will be issued for online DBRC regions and processing will end for non-online DBRC regions.

## Records in the RECON data sets

The RECON data set contains many types of records. Some records, such as header records, exist primarily to control processing of the RECON data set. Other records exist to define the various data sets used in the recovery of DBDSs. Still others exist to record events related to the use of DBDSs.

The following figure shows the major relationships between RECON record types. These relationships are described in detail in the topics that follow the figure.






Figure 55. Major RECON record types and the relationships between them

#### Related reading:

- See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for details of what the records contain.
- See “Database Recovery Control service aids” in *IMS Version 11 Diagnosis* for information about RECON keys.

## RECON header records

The RECON data set has two record types that contain control information DBRC uses in its processing.

- RECON header

In addition to internal control information that DBRC creates and uses, the RECON header record contains information provided by you in the parameters of the INIT.RECON and CHANGE.RECON commands.

- RECON header extension

The RECON header extension record contains RECON data set configuration and state data that DBRC uses to process the RECON data sets.

## Log data set records

Three types of log data sets exist. Each type can have a primary log record, secondary log record, interim-primary log record, and interim-secondary log record.

The log data set records for each log data set are:

- Recovery log data set

- PRILOG (primary log record)
- SECLOG (secondary log record)
- IPRI (interim primary log record)
- ISEC (interim secondary log record)
- System log data set
  - PRISLD (primary system log data set)
  - SECSLD (secondary system log data set)
  - PRITSLDS (primary RSR tracking site log data set)
  - SECTSLDS (secondary RSR tracking site log data set)
  - IPRISL (interim primary system log data set)
  - ISECSL (interim secondary system log data set)
  - IPRITSLD (interim primary RSR tracking site log data set)
  - ISECTSLD (interim secondary RSR tracking site log data set)
- Online log data set
  - PRIOLDS (primary online log data set)
  - SECOLDS (secondary online log data set)
  - IPRIOL (interim primary online log data set)
  - ISECOL (interim secondary online log data set)

Log records come in sets called *PRILOG families*. A PRILOG family consists of a PRILOG and one or more of the following: SECLOG, PRISLD, and SECSLD for a given time period and IMS subsystem. All records in this set have the same start and end times and normally have matching data set entries. The same LOGALL record applies to all members of the set.

DBRC creates the PRILOG and PRISLD records whenever an online IMS opens the first OLDS, and updates them each time an OLDS is archived. If you use dual archiving, DBRC creates SECLOG and SECSLD records when the first OLDS is archived and updates them each time an OLDS is archived.

**Related reading:** See “Archiving an OLDS” on page 486 for more information about DBRC and the archiving of online logs.

Log data sets output from IMS batch jobs are recorded in PRILOG / SECLOG records even though, technically, they are SLDSs. These records are created whenever the output log is opened and updated when volume switches occur.

In addition, during log recovery processing, DBRC creates an IPRISL or IPRIOL record for each interim primary-log data set and an interim secondary-log record for each interim secondary-log data set whenever the Log Recovery utility runs. An interim log record is an internal record that is used to reflect intermediate processing of the DUP function of the Log Recovery utility.

**Related reading:** See “Log Recovery utility (DFSULTR0)” in *IMS Version 11 System Utilities* for more information about interim primary and interim secondary-log data sets.

Although not part of normal DBRC operation, you can use the following commands to create log records (for example, to set up a test environment or for RECON data set repair purposes):

- NOTIFY.PRILOG

- NOTIFY.SECLOG

## Database recovery records

In addition to the log data set records, the RECON data set contains records with database recovery information in them.

**Related reading:** See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for a listing of the records in the RECON data set and tables that describe the fields in the RECON data set records.

### Backout record (BACKOUT)

A BACKOUT record contains information about units of recovery including time stamp, associated PSB name, recovery token, and database name. The names of nonrecoverable databases are not stored in the BACKOUT record.

### Change accumulation group record (CAGRP)

A CAGRP record identifies a change accumulation (CA) group. This record includes up to 32,767 names of DBDSs whose change records are accumulated during one run of the Database Change Accumulation utility. Each DBDS (for which DBRC is controlling recovery) can be a member of only one CA group in order for the Database Change Accumulation utility to accumulate its changes. You specify the CAGRP name in the INIT.CAGRP command that you use to create a CAGRP record in the RECON data set.

The CAGRP record contains the name of a member of a partitioned data set. This member contains the skeletal JCL that is used to generate the JCL to run the Database Change Accumulation utility for this CA group. The CAGRP record also contains an indicator that specifies whether change accumulation data sets that correspond to this group can be reused. It contains an indication of the maximum number of change accumulation data sets that are maintained for the group.

### Change accumulation run record (CA)

A CA record contains information about a change accumulation data set and can be either available or in use. For each CAGRP record, there can be up to 1024 CA records.

An available CA record is created by an INIT.CA command. This CA record contains the volume serial numbers and the data set name of a data set that is to be used for output from a subsequent run of the Database Change Accumulation utility. You can create available CA records only for those CA groups that are defined with the REUSE parameter.

An in-use change accumulation record is created by a run of the Database Change Accumulation utility. It can be a formerly available CA record that was used during a run of the Database Change Accumulation utility, or it can be a new record with information obtained from the JCL. The information in an in-use change accumulation record includes:

- Data set name
- Volume serial numbers
- Run time of the Change Accumulation utility
- Stop time of the last log volume that the Database Change Accumulation utility processed; or, if the CA processed a subset of logs, the start time of the first log that should be processed on the next execution

You can use a NOTIFY.CA command to create a CA record.

## Data group record (DBGRP, DBDSGRP, RECOVGRP)

You can use a DBDSGRP record to define the following types of named groups: DBDSGRP, DBGRP, and RECOVGRP.

The named groups are:

- DBDSGRP, a group of DBDSs and DEDB areas
- DBGRP, a group of DL/I DBs and DEDB areas that can be named in the DATAGROUP parameter for the /STA, /STO, and /DBR commands
- RECOVGRP, a group of DL/I DBs and DEDB areas that are logically related for database recovery purposes

All groups must have unique names. For example, a DBDSGRP cannot have the same name as a DBGRP.

You use the INIT.DBDSGRP, CHANGE.DBDSGRP, DELETE.DBDSGRP, and LIST.DBDSGRP commands to manipulate all three data group types.

**Recommendation:** Although any type of group can be named in the DATAGROUP parameter for the /STA, /STO, and /DBR commands, the use of a DBDSGRP is not recommended because it is inefficient.

**Related reading:** See “DBDS group considerations” on page 529 for more information on DBDS groups.

## Database records (DB)

DBRC treats DL/I, Fast Path, and HALDB (High Availability Large Database) database records differently.

These types of records, their treatment, and contents are explained as follows:

### DL/I database records

A DB record identifies a database that is registered and whose recovery is under the control of DBRC. This record contains information about the database and related recovery information including:

- Database name
- Database type
- Share level of database
- List of subsystems using the database
- Extended Error Queue Element (EEQE) counter
- IRLM identification of the first subsystem that authorized the database (if IRLM is used)

A DBDS record identifies a DBDS whose recovery DBRC is to control. This record contains information about the DBDS (such as its data set organization) and related recovery information including:

- Name of the CA group to which the DBDS belongs
- Maximum number of image copy data sets to be maintained for this DBDS
- Indication of whether image copy data sets are to be reused
- Period of time that image copy data sets are to be maintained for this DBDS
- Name of the implied skeletal JCL default member
- Extended Error Queue Elements (EEQEs)

- Names of the members of the partitioned data set of skeletal JCL to be used in order to generate JCL for utilities that are run for this DBDS

To describe DL/I databases and DBDSs, DBRC maintains logically related records, as shown in the following figure.




Figure 56. DBRC DL/I records

### HALDB records

From a RECON data set perspective, HALDBs consist of a HALDB master (TYPE=HALDB) and one or more HALDB partitions (TYPE=PART).

The following figure shows the relationship of the RECON data set records that represent a HALDB.




Figure 57. Record structure for a HALDB in the RECON data sets

*HALDB Master (DSPDBHRC)*

The RECON data set stores information pertaining to the entire HALDB by using a DB header record. The DB header record includes the following:

- HALDB master name
- Global DMB number
- Current change version number
- TYPE=HALDB
- Data sharing level
- RSR Global Service Group Name and tracking level
- Recovery group
- HALDB Partition Selection exit routine name
- Database and data set organization
- Partition ID of the latest partition defined
- Recoverability of the HALDB
- Number of partitions in the HALDB
- Capability for online reorganization
- Number of data set group members

The TYPE=HALDB DB record stores information about the HALDB. Applications conduct DB activity at the HALDB master level; DBRC conducts the DB activity at the HALDB partition level. A subsystem authorizes a HALDB partition, not a HALDB master.

#### *HALDB Partition*

Each partition of the HALDB consists of the following RECON data set records:

- Partition record (DSPPTNRC)  
DSPPTNRC contains information that applies to the individual partition. The HALDB Partition Definition utility displays the partition record information. The LIST command does not display the partition record.
- Partition DB record (DSPDBHRC)  
DSPDBHRC accesses the HALDB at the partition level. Like the TYPE=IMS DB record, the DB record for the HALDB partition records all sharing and recovery information. The partition name sets the database name field in this record. TYPE=PART has been defined for this record. The following fields have the same settings for each partition across the entire HALDB:
  - Global DMB number
  - Data sharing level
  - RSR Global Service Group Name and tracking level
  - Recoverability of the partitions
  - HALDB master name
  - Capability for online reorganization
- Partition DBDS records (DSPDSHRC)  
Depending on the organization of the HALDB, there can be three types of DBDSs for each HALDB partition: data, index, and ILDS data sets. Multiple data DBDSs can exist, but only one of each of the others. Only data DBDSs can be recovered. The other DBDSs are rebuilt using the HALDB Index/ILDS Rebuild utility (DFSPREC0).

**Related reading:** See *IMS Version 11 Database Administration* for information about the data set and DDN naming conventions for DBDS records.

### Fast Path database records

To describe DEDBs, areas, and Area Data Sets (ADSs), DBRC has a logical structure of records, as shown in the following figure.




Figure 58. DBRC Fast Path database records

DBRC uses the DB and DBDS records to describe both DL/I databases and DEDBs; however, DBRC adds an ADS list to the Fast Path DBDS record giving information about each ADS. Each DEDB may contain multiple areas, and each area may contain up to seven ADSs.

The Fast Path DB record contains information similar to the information in a DL/I DB record, except that it describes a DEDB; and it does not contain a list of authorized subsystems. For Fast Path, this list is in the DBDS record, which is composed of an area authorization record and an area recovery record. The Fast Path DB record is displayed in the listing as the DBDS record.

**Related reading:** See the sample RECON listing in *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*.

When an area is registered in the RECON data set, it ensures that:

- The area names in the DEDB are unique
- The ADS DD names in an area are unique
- No more than seven ADSs are defined for an area

### Global service group record (GSG)

A GSG record defines a global service group and the service groups that make up the GSG. An RSR service group is made up of two service groups: the active and the tracker.

GSG records are created by the INIT.GSG command and can be deleted by the DELETE.GSG command. The INIT.SG and DELETE.SG commands add and remove service group definitions to and from the GSG record. The CHANGE.SG command modifies information about a service group.

## Image copy record (IMAGE)

An IMAGE record contains information about an image copy data set and can be either *available* or *in-use*.

An available IMAGE record is created by an INIT.IC command. It describes data sets that are to be used for output from a subsequent run of an image copy utility. You can create available IMAGE records only for those DBDSs that are defined with the REUSE parameter.

An in-use IMAGE record is created by an execution of an image copy utility:

- If the DBDS or area is defined with REUSE, it is a formerly available IMAGE record or a reused in-use record.
- If the DBDS or area is defined with NOREUSE, it is a new record with the data set description that is obtained from the JCL.

In addition to the data set description, the in-use IMAGE record identifies the type of copy operation and contains the copy operation's run time and, depending on the type of copy, the copy operation's stop time.

If you request that the image copy utility create an image copy data set and a duplicate image copy data set, DBRC records information about both in the same image copy record. The first record is designated IC1, and the duplicate is designated IC2.

If you create a nonstandard image copy data set (one that the image copy utility did not create), you must use a NOTIFY.UIC command to record its existence in the RECON data set.

**Related Command:** See “Image copy utilities (DFSUDMP0, DFSUDMT0, DFSUICP0)” on page 516 for more information on the IMAGE COPY record.

## Reorganization record (REORG)

DBRC creates a REORG record each time you reorganize a data set of a registered database or partition by using one of the following methods.

- HALDB Online Reorganization
- The HISAM Reorganization Reload utility
- The HD Reorganization Reload utility

For a HALDB Online Reorganization, the REORG record indicates that the type of REORG is online, includes a stop time, and indicates if the reorganized data set can be used for recovery. DBRC creates a REORG record in the RECON data set for each database data set that you reorganize.

**Related concepts:**

 HALDB online reorganization (Database Administration)

**Related reference:**

 HD Reorganization Reload utility (DFSURGL0) (Database Utilities)

 HISAM Reorganization Reload utility (DFSURRL0) (Database Utilities)

## Log allocation record (LOGALL)

A LOGALL record identifies the registered DBDSs that were changed while the corresponding log data set was open.



DBRC creates a log allocation (LOGALL) record for each PRILOG record.

### **Database allocation record (ALLOC)**

For Fast Path, DBRC creates an ALLOC record when the area is placed in OPEN-for-update status. For DL/I, DBRC creates an ALLOC record when IMS updates a DBDS the first time during a run of IMS or the first time after you enter a /DBRECOVERY command. This record contains the time stamp of its creation and the time stamp of the opening of the corresponding log. This time stamp identifies the log data sets containing the database change records that are needed for recovery. If the DBDS is subsequently deallocated, DBRC adds the time stamp of the deallocation to the ALLOC record.

DBRC automatically deletes allocation records if they are older than the oldest image copy record and if DBRC no longer needs them for database recovery. As DBRC deletes an ALLOC record, it changes the associated LOGALL record. This is part of the DBRC Image Copy utility exit routine processing. This automatic deletion of ALLOC records for a DBDS does not occur under either of the following conditions:

- The ALLOC record has no deallocation time.  
A deallocation time is recorded when the database or area has had the /DBR command run against it. Otherwise, the ALLOC record uses the log close time as an implicit deallocation time.
- The PRILOG record associated with the ALLOC record is open (it has a STOPTIME of zero) but is not marked in error.

In cases such as these, you can use the DELETE.ALLOC command to delete unwanted ALLOC records from the RECON data set. DBRC automatically updates allocation time stamps during online and concurrent image copy utility exit routine processing to move the allocation time stamps forward in time.

### **Recovery record (RECOV)**

DBRC creates a RECOV record each time you run the Database Recovery utility to recover a DBDS.

The RECOV record can indicate one of these types of recovery:

- A full recovery of a DBDS where the RECOV record contains the time stamp of the recovery.
- A time-stamp recovery where the RECOV record contains the time stamp of the run of the Database Recovery utility and the time stamp to which the DBDS was recovered.

### **Subsystem records (SSYS)**

The RECON data set uses the subsystem (SSYS) record to describe data sharing information.

An SSYS record is created when an IMS subsystem signs on to DBRC. This SSYS record contains information about the subsystem and related recovery information including:

- Subsystem name and type (online or batch)
- IRLM identification
- Abnormal-end flag and the recovery-process start flag
- List of authorized databases
- Time stamp that correlates the subsystem entry with the appropriate log records

---

## Mapping the records in the RECON data set

To find the DSECTS defining the formats of the RECON records in DBRC in ADPSMAC and SDFSMAC, run the generate job with SDFSMAC=ALL.

The member names are:

**Member (DSECT) name**

**RECON record that the DSECT maps**

**DSPALLRC**

ALLOC record

**DSPBKORC**

BACKOUT record

**DSPCAGRC**

Change accumulation group (CAGRP) record

**DSPCHGRC**

Change accumulation (CA) record

**DSPDBHRC**

Database (DB) record

**DSPDGRC**

DBDSGRP record

**DSPDSHRC**

DBDS record

**DSPGSGRC**

Global Service Group record

**DSPIMGRC**

Image copy (IC) record

**DSPLGARC**

LOGALL record

**DSPLOGRC**

PRILOG / SECLOG / PRISLD / SECSLD / PRITSLDS / SECTSLDS record

**DSPOLDRC**

PRIOLD / SECOLD record

**DSPPTNRC**

HALDB partition record

**DSPRCNRC**

HEADER record

**DSPRCR1**

HEADEREXT record

**DSPRCVRC**

RECOV record

**DSPRRGRC**

REORG record

**DSPRDTRC**

RECON DMB table record

**DSPSSRC**

SUBSYS record

---

## Maintaining the RECON data sets

You should perform periodic maintenance on the RECON data set to maintain data integrity and an acceptable level of performance.

The following topics discuss the main tasks involved.

### Backing up the RECON data sets

Back up the RECON data sets frequently. They are a critical resource. Always make backup copies of the RECON data set after performing any RECON record maintenance, such as registering databases and adding or deleting change accumulation groups.

Use the BACKUP.RECON command to perform backup. This command issues the necessary RESERVE commands (reserving the device during backup processing) to ensure backup integrity. Then it invokes the AMS REPRO command to copy the data set. The BACKUP.RECON command copies only the first copy of the RECON data set. Its parameters determine whether it makes one or two copies. In parallel access mode, no reserve commands are performed. Instead RECON activity is quiesced for all active DBRC instances.

If parallel access is enabled, the user has the option to take a backup copy of a Transactional VSAM (TVS) RECON using DFSMSdss COPY or DUMP. When DFSMSdss COPY or DUMP is used, TVS automatically quiesces the data set so that a consistent or “sharp” backup copy is produced.

**Related reading:** See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for additional information on the BACKUP.RECON command.

### Deleting unnecessary RECON records

You can delete unnecessary RECON records using the following methods.

**Related concepts:**

“Recovery period of image copy data sets and GENMAX” on page 521

#### Automatic deletion of extraneous records

Normally you do not need to perform much record maintenance for database-related records.

When the RECON data set is notified of an image copy, it may delete or reuse the oldest in-use IMAGE record, and a later IMAGE record becomes the oldest IMAGE record. RECOV and REORG records with start times earlier than the (now) oldest IMAGE record, and ALLOC records with DEALLOC times earlier than the earliest time, are now extraneous, and are deleted from the RECON data set. This is the image copy cleanup process.

When extraneous IMAGE records are deleted from the RECON data set, all active ALLOC records are updated to the time of the first log volume necessary for recovery, based on the oldest image copy for the DBDS or area. When the cleanup process deletes an extraneous ALLOC record, it changes the state of the associated LOGALL records. After all the ALLOC records associated with the LOGALL record have been deleted (this may take place over many image copies for many databases), the PRILOG record associated with the LOGALL record becomes inactive.

## PRILOG compression

PRILOG record compression is the deletion of all inactive data set entries in an open PRILOG record.

A data set entry is defined as being inactive when it is older than all of the following criteria:

- Log retention period
- Oldest allocation (ALLOC) for any database updated on that log
- Earliest restart checkpoint for the online IMS

PRILOG record compression deletes inactive data set entries up to the oldest ALLOC on the log or the first gap in the data set entries. A gap occurs when an OLDS has not yet been archived.

When inactive data set entries are deleted from active PRILOGs, they are compressed to a single dummy data set entry that has the same start time as the start time of the log and the same stop time as the stop time of the last inactive data set entry deleted.

PRILOG compression can be performed using the following two methods:

### Automatic PRILOG compression

PRILOG record compression is attempted automatically after an OLDS is archived. At an RSR (Remote Site Recovery) tracking site, automatic compression is attempted when a tracking log data set is opened by the log router and recorded in the RECON data set.

Compression includes the deletion of all inactive data set entries in the PRILOG record. When applicable, corresponding entries in the SECLOG, PRISLD, and SECSLD records are also deleted.

### Manual process for PRILOG compression

You can initiate PRILOG record compression manually by using the DELETE.LOG INACTIVE command. This command deletes inactive data set entries from active PRILOG records and deletes entire inactive PRILOG records.

If your PRILOG is not getting compressed, you will receive message DSP1150I for one of the following reasons:

- EARLIEST CHECKPOINT

If this time stamp is within the first DSN entry in the PRILOG, then compression is not possible. Most often, the oldest checkpoint needed for restart causing the problem is the checkpoint to rebuild the message queues. If so, take a SNAPQ checkpoint. When the OLDS that records this SNAPQ checkpoint is archived, the earliest checkpoint time stamp is updated.

- EARLIEST ALLOC TIME

The databases listed in the LOGALL record are in earliest allocation time stamp order. The earliest allocation time listed in the LOGALL record is associated with the first databases in the listing.

Concurrent image copies (CICs) and Online image copies (OICs) update the allocation time to reflect the earliest log that is needed for recovery. The checkpoint IDs and checkpoint counts on the logs before the start of the CIC are

used to determine where the new allocation time is set. If the allocation times are not being moved forward-in-time, ensure that checkpoint IDs and counts are being recorded in the DSN entries.

Old allocations for any databases in the LOGALL record can cause compression to fail. Image copy databases on a regular basis to allow old allocations to be deleted.

- LOG RETENTION TIME

Indicates the minimum amount of time that DBRC is to keep log records in the RECON data set.

## Manual log record deletion

You can manually delete log records by using the DELETE.LOG INACTIVE or the CLEANUP.RECON command.

DBRC does not automatically delete RECON records that describe log data sets (PRILOG and SECLOG records). This design gives you control over which RECON records associated with log data sets are deleted. You must periodically delete PRILOG and SECLOG records that are no longer needed for recovery.

The current size of the PRILOG record is printed on the listing of the PRILOG record.

Use the DELETE.LOG INACTIVE command to delete inactive PRILOG records, SECLOG records, and their associated LOGALL records.

Common problems are old open (nonzero stop time) PRILOG records that were created by jobs that terminated abnormally (without closing their logs). You can find these open logs by issuing the LIST.LOG command with the OPEN parameter. You can use the DELETE.LOG command with the STARTIME parameter to remove old, unnecessary, open PRILOG family records.

Alternatively, use the CLEANUP.RECON command to unconditionally delete old or expired log information and recovery-related information from the RECON data set. Recovery-related information includes image copy, allocation, reorganization, and recovery records. Using the specified time stamp or retention period, recovery-related data and inactive logs are deleted and PRILOG compression is forced on both open logs and closed logs which have a close time after this time.

Deleting log records does not prevent the RECON data sets from filling up because the space freed by deletions might not be reused by VSAM. However, if you delete the log records before you back up or reorganize the RECON data set, you can reclaim the space during backup or reorganization.

**Attention:** To avoid deleting necessary recovery and log information from the RECON data set, run the CLEANUP.RECON command on a copy of the RECON data set first and verify the results, which indicate the length of time that is needed by the job. You can then use this information to determine when to run this command on your active RECON data set.

**Related reference:**

 CLEANUP.RECON command (Commands)

 DELETE.LOG command (for OLDS) (Commands)

 DELETE.LOG command (for RLDS and SLDS) (Commands)

 LIST.LOG command (for a PRILOG family) (Commands)

## Reorganizing the RECON data sets

You need to reorganize the RECON data sets periodically. Many of the record keys in the RECON data set include the date and time. DBRC recording of IMS log and database activity can cause CI and CA splits, that can degrade performance.

In addition, deleting unnecessary records might not prevent the RECON data set from filling up because VSAM does not always reuse the space freed.

### Using CHANGE.RECON to reorganize a RECON data set

You can reorganize a RECON data set online if you are using dynamic allocation for the RECON data set by using the CHANGE.RECON command. A spare RECON data set must be available.

In this situation, you can issue the CHANGE.RECON command with the REPLACE option. This causes DBRC to copy and reorganize the active RECON data set (specified on the CHANGE command) to the spare data set. VSAM removes all CI and CA splits, and restores the original FREESPACE attributes.

The CHANGE command also deallocates the old RECON data set (the one that needed reorganization). Before you can delete and redefine this data set, however, you must wait for all other subsystems that are using it to deallocate it. If you redefine the data set with the same name it originally had, it is available to the online system for use as a spare data set. You can repeat this process to reorganize the second active RECON data set. If you use JCL in order to allocate data sets, dynamic deallocation does not occur.

If you do not use dynamic allocation or if a spare RECON data set is not available, you must wait for the online subsystem and all other subsystems that access the RECON data set to deallocate it before you can reorganize.

### Backing up RECON data sets before reorganization

Back up the RECON data sets before and after reorganizing, using the following procedure.

1. Copy the RECON data sets to temporary data sets with different data set names.
2. Verify that the copied data sets are uncorrupted.
3. Delete the original data sets.
4. Redefine the original data sets using the same data set names.
5. Copy the temporary data sets back to these original data sets.

### Procedure for reorganizing the RECON data sets

To clean up your logs and reorganize your RECON data sets, follow this procedure. This process assumes that you are running DBRC in dual mode with RECON1 (as Copy1), RECON2 (as Copy2), and a spare.

1. Issue a LIST.LOG OPEN command to identify any open logs. Determine which of the identified logs should be open, and which should be closed or deleted. Close any logs that should be closed or deleted.
2. Issue a DELETE.LOG INACTIVE command to delete any inactive or unused logs. You are now ready for the RECON data set reorganization.
3. Issue a CHANGE.RECON with the REPLACE option for RECON1. DBRC then:

- a. Signals the start of the RECON data set reconfiguration with message DSP0380I
  - b. Identifies the subsystems that are active at reconfiguration with message DSP0388I
  - c. Signals that the process has completed with message DSP0381I.
4. Replace the discarded RECON data set. See “Replacing a discarded RECON data set” on page 561 for more information.

Your RECON data sets have now been reorganized and the logs are now clean.

## Replacing damaged RECON data sets

If an I/O error occurs in the current RECON data set that is the only one remaining, DBRC stops the job. Any other jobs that are currently using the RECON data set continue to run if no other I/O error is encountered.

If an I/O error occurs in a RECON data set and two RECON data sets exist, DBRC attempts to locate a spare data set. If a spare is available, DBRC copies the RECON data set without the I/O error to the spare RECON data set. DBRC then establishes the spare as the Copy2 RECON data set.

**Recommendation:** After the spare RECON data set replaces the RECON data set that had the error, redefine the discarded RECON data set as quickly as possible. If you immediately replace the RECON with the I/O error, you are unlikely to experience a subsystem failure due to loss of all RECON data sets. See “Replacing a discarded RECON data set” on page 561 for more information.

If DBRC cannot locate a spare RECON data set and you have specified the STARTNEW parameter of the INIT.RECON command, DBRC continues processing with one RECON data set. Otherwise, DBRC completes the current job but does not start new jobs until you define a spare RECON data set.

## Recovering the RECON data sets

Steps for recovering your RECON data set depend on the set up of your RECON data set and the situation.

The following paragraphs discuss various scenarios for recovering a RECON data set.

### A Spare RECON data set is available

If an I/O error occurs on a RECON data set and a spare data set is available, DBRC copies the good RECON data set to the spare, and then activates the spare.

If, however, you want to analyze the RECON data set error, before deleting and redefining the discarded RECON data set, make a copy of it for later problem diagnosis.

### A Spare RECON data set is not available

If a spare RECON data set is not available, all currently executing jobs continue processing using the RECON data set in single mode. If you specified the STARTNEW parameter in the INIT.RECON or CHANGE.RECON command, DBRC allows new jobs to start with only one RECON data set. This is not recommended as it jeopardizes the integrity of the system.



If one of the data sets in the set of RECON data sets becomes unusable by DBRC, you need to deallocate the RECON data set that is unusable and allocate a new spare.

In an RSR environment, the isolated log sender (ILS) starts its own copy of DBRC. If Automatic RECON Loss Notification is not active, you might need to stop ILS to terminate the DBRC in the Transport Manager Subsystem (TMS) address space. This causes the ILS's DBRC to deallocate the RECON data sets so that you can replace the unusable RECON data set. Issue a STOP ILS(gsg) command for each started ILS instance. Then issue START ILS(gsg) to bring up ILS and DBRC again.

**Both RECON data sets are not usable:**

It is unlikely that both RECON data sets would be not usable. If, however, both RECON data sets ever become unusable, follow this procedure:

1. Stop all jobs that require access to the RECON data set.
2. Optional: Use the AMS REPRO command to back up the RECON data sets if you can access both of them.
3. Use the AMS utility to delete and redefine your RECON data sets.
4. Use the AMS REPRO command to restore one of the RECON data sets.
5. Use the AMS REPRO command to restore the other RECON data set from the first.
6. Use the LIST.RECON command to list one of the RECON data sets. Evaluate the list and determine which DBDSs have been updated since you made the backup in step 2. If you cannot determine which DBDSs have been updated, assume that all have been updated.
7. Use the CHANGE.IC command with the INVALID parameter to mark all image copy records that are in error for all applicable DBDSs in step 6.
8. Make an image copy of all applicable DBDSs in step 6.
9. Use the BACKUP.RECON command to make a backup copy of the RECON data sets.

The RECON data sets are now restored and resynchronized with the databases.

If you do not control an excessive number of databases, it might be easier to follow this procedure:

1. Stop all jobs that require access to the RECON data set.
2. Define new RECON data sets.
3. Initialize these RECON data sets.
4. Register the environment. Always keep a backup copy of the most recently initialized, but not-yet-used, RECON data set available.
5. Take image copies of all databases.

Finally, before you proceed with your regular operations, clean up the new RECON data set by, for example, closing any open, out-of-date OLDSs with the NOTIFY.PRILOG command.

## **Replacing a discarded RECON data set**

DBRC detects that a RECON data set is discarded only when some activity occurs that causes DBRC to access the RECON data sets. You can have multiple instances of DBRC whenever you have multiple IMS subsystems, online or batch. You cannot delete and redefine a discarded RECON data set until all instances of DBRC detect that a change has occurred and they deallocate the discarded data set.



DBRC lists the subsystems that are active at reconfiguration in message DSP0388I. This message enables you to identify the subsystems that might need your help in detecting the status change of the RECON data sets.

To redefine a RECON data set after an I/O error has occurred, or in conjunction with the CHANGE.RECON REPLACE command, follow this procedure:

1. Allow all batch jobs using DBRC to finish.
2. Issue LIST.RECON STATUS in all online subsystems if you do not have Automatic Loss Notification active. Issuing the command causes the online subsystems to obtain the same Copy1 and Copy2 RECON data sets and to deallocate the discarded RECON data set. If you do have Automatic Loss Notification active, all subsystems are automatically notified to deallocate the discarded RECON data set. See "RECON loss notification" for more information on RECON Loss Notification.
3. Use the AMS DELETE command to delete the discarded RECON data set.
4. Use the AMS DEFINE command to recreate the RECON data set as an empty VSAM KSDS data set. Use the same procedure that you used originally to create the RECON data set. See "Creating a RECON data set" on page 542.

## RECON loss notification

All DBRC instances that are allocated to the RECON data set, using the same IMSplex, are notified of errors, at the time the errors occur, through the IMS Structured Call Interface (SCI). Then, any DBRC allocated to the discarded RECON data set will deallocate.

In parallel RECON access, DBRC must ensure that no other DBRC instance can access the RECON data set during the reconfiguration process. At that point, all the other DBRCs re-access in parallel mode.

DBRC obtains the IMSplex name from the DBRC SCI Registration exit routine (DSPSCIX0), or from an EXEC parameter, IMSPLEX. Any attempt at accessing a RECON data set using a different IMSplex, or no IMSplex, is rejected and message DSP1136A is issued. Otherwise, RECON Loss Notification is automatically invoked.

If you want to change the IMSplex associated with a set of RECONs, use the IMSPLEX or NOPLEX parameter on the CHANGE.RECON command. To change an IMSplex-to-RECON association when a DBRC instance is active:

1. Wait for all DBRC activity on the current IMSplex to stop.
2. Submit a Database Recovery Control utility job to change the IMSplex name.  
If parallel access is enabled and you want to remove the IMSplex name from the RECON data set by using the CHANGE.RECON NOPLEX command, parallel access must be disabled first by using the CHANGE.RECON ACCESS(SERIAL) command.
3. Alter the IMSplex name in the DSPSCIX0 exit routine.
4. Ensure that the new IMSplex SCI is ready.

### **Related reading:**

- See "IMSplex overview" on page 19 for introductory information about IMSplexes.
- See *IMS Version 11 Exit Routines* for more information on the DBRC SCI Registration exit routine (DSPSCIX0).

- See *IMS Version 11 System Definition* for more information about the IMSPLEX EXEC parameter.

## DBRC security

The RECON data sets are critical to the integrity of IMS databases. Therefore, consider restricting access and providing access to only a subset of the DBRC commands for those users who must issue them.

In the online region, IMS provides some authorization functions for commands issued. By default, IMS restricts /RMxxxxxx commands (except for /RMLIST) to the Master Terminal Operator (MTO).

You can establish authorization control for DBRC commands through RACF (or an equivalent security product), a user exit routine, or both. For example, logon ID SMITH could have access to issue an INIT.DB command against DBD PAYROLL but not against DBD CUSTOMER.

The HALDB Partition Definition utility is an ISPF application that allows you to manage the definitions of IMS HALDBs and their partitions in the RECON data set, providing functionality equivalent to the following:

- INIT.DB
- INIT.PART
- CHANGE.DB
- CHANGE.PART
- CHANGE.DBDS
- DELETE.DB
- DELETE.PART
- LIST.DB

Authorization for the HALDB Partition Definition utility is controlled using the same resources as defined for these commands with one exception: the CHANGE.PART resource is used in place of the CHANGE.DBDS resource.

**Related reading:** See *IMS Version 11 Database Utilities* for more information about the HALDB Partition Definition utility.

### Authorizing DBRC commands and API requests with RACF

You can set up DBRC command and API request authorization support with RACF by defining resource profiles (explicit or generic) that cover all of the DBRC commands and API requests, and permitting appropriate user access to these profiles.

The resource name corresponds directly to the command and consists of a high-level qualifier (HLQ) and up to three elements of the DBRC command. The resource name model has the following format:

- High-level qualifier (for example, SAFHLQ)
- Command verb (for example NOTIFY)
- Command modifier (for example, PRILOG or IC)
- Command qualifier (for example, dbname or OLDS)

DBRC uses a list of resource names for command authorization support. See Chapter 43, “Resource names for command authorization,” on page 579.

Different profiles can be used based on the set of RECON data sets being used, but only one HLQ name per RECON data set is allowed. For example, you might use an HLQ name of PRODRECN for the production RECON data sets and an HLQ name of TESTRECN for the test RECON data sets.

Other considerations for your definitions are as follows:

- Resource profiles that protect the DBRC commands must be defined in the FACILITY resource class.
- User identifiers in the ACEEUSRI field of the accessor environment element (ACEE) are used for all authorization checking and the user has READ access to the resource profile, at a minimum.
- DBRC command authorization is enabled or disabled by using the CHANGE.RECON and INIT.RECON command with the CMDAUTH keyword. The HLQ name is required when enabling command authorization.

### **Authorizing security profiles for DBRC commands and API requests**

You can define security resource profiles for your system to protect specific DBRC commands and requests, and permit appropriate user access to each command or request.

The security resource name corresponds directly to the DBRC command or request and consists of a high-level qualifier (HLQ) and up to three elements of the command or request. Assume that you are using RACF for data security and decide to use IMSA as your HLQ to identify your IMS system.

To restrict access to the DELETE.SUBSYS command:

1. Issue the following RACF command to allow access to all DBRC commands:  
`RDEFINE FACILITY IMSA.** OWNER(...) UACC(READ)`
2. Issue the following two RACF commands to restrict the DELETE.SUBSYS command to the user IDs listed in the ID parameter:  
`RDEFINE FACILITY IMSA.DELETE.SUBSYS.* OWNER(...) UACC(NONE)`  
`PERMIT IMSA.DELETE.SUBSYS.* CLASS(FACILITY) ID(...) ACC(READ)`  
Command protection is now in effect.
3. Issue the following command to clear the cache and cause RACF to refresh and read the new values from the database:  
`SETROPTS REFRESH GENERIC(FACILITY) RACLIST(FACILITY)`
4. Issue the following DBRC command to specify accesses to the RECON data set. For example, to use the prefix IMSA for security checking:  
`CHANGE.RECON CMDAUTH(SAF,IMSA)`

### **Authorizing DBRC commands and API requests with the DSPDCAX0 exit routine**

You can use the DBRC Command Authorization exit routine (DSPDCAX0) to verify that a user is authorized to issue a particular command or API request.

If you use this exit, you will make the appropriate selection when using the INIT.RECON or CHANGE.RECON commands. Refer to *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for more information.

The DSPDCAX0 exit routine is required if it is specified to authorize DBRC commands (that is, if either sub-parameter EXIT or BOTH is specified in the CMDAUTH keyword on the INIT.RECON or CHANGE.RECON commands). The

DSPDCAX0 exit routine must be specified in an authorized library or in LINKLST. If the DSPDCAX0 exit routine is in a concatenated STEPLIB or JOBLIB, only the data set that contains the DSPDCAX0 exit routine must be authorized. If the DSPDCAX0 exit routine is in LINKLST, no authorization check is performed.

**Related reading:** See *IMS Version 11 Exit Routines* for more information on the DBRC Command Authorization exit (DSPDCAX0).

### **Authorizing DBRC commands and API requests with both the DSPDCAX0 exit routine and RACF**

The DBRC command authorization exit routine (DSPDCAX0) can be used in conjunction with RACF to provide command authorization.

In such cases, the security product is invoked first. The return and reason codes are passed to the exit routine. The return code that the exit routine issues ultimately determines the success or failure of the command authorization; the exit routine overrides the outcome of the security product. Therefore, DBRC messages that are a result of unsuccessfully invoking the security product will be suppressed.

Again, if EXIT or BOTH is specified in the CMDAUTH keyword on the INIT.RECON or CHANGE.RECON commands, DSPDCAX0 is a required exit routine. Also, DSPDCAX0 must be found in an authorized library or in LINKLST. If DSPDCAX0 is found in a concatenated STEPLIB or JOBLIB, only the data set containing DSPDCAX0 must be authorized. If DSPDCAX0 is found in LINKLST, no authorization check is performed.

**Related reading:** See *IMS Version 11 Exit Routines* for more information about the DBRC Command Authorization exit (DSPDCAX0) routine.

### **Overriding DBRC security on the RECON data set**

You can create a copy of the production RECON data set for test purposes or to send to IBM Software Support to help re-create a problem with the original RECON data set.

A copy of the RECON data set inherits the security settings of the original RECON data set. Disable DBRC security for a copy of the RECON data set when the copy needs to be accessed by someone lacking the level of authorization required by the current security settings.

To create a copy of the RECON data set that does not inherit the security settings of the original RECON data set:

1. Issue the CHANGE.RECON command with the *rcnqual* option. If the *rcnqual* value is not a substring of the RECON COPY1 data set name, the command bypasses security. However, if the *rcnqual* value is a substring of the RECON COPY data set name, the command bypasses security only if the *rcnqual* value is followed by an asterisk (\*) and not at the beginning of the data set name. For example, if *rcnqual*=IMSTEST\*, and IMSTESTL.RECON1 is copied into RCNCOPY.IMSTESTL.RECON1, security would be bypassed because the copied data set name contains the *rcnqual* value, but does not begin with it. If *rcnqual*=IMSTEST, security would not be bypassed for the copied data set.
2. Create the copies of the production RECON data set with names that do not contain the string that was specified in the *rcnqual* option. These copies of the RECON data set can be accessed by those without command authorization.

| When an attempt to use the copies is made, the system compares the string value  
| stored in the RECON data set against the RECON COPY1 data set name. Once it  
| verifies that the strings do not match and the string in the copy is not a substring  
| of the RECON data set, message DSP1211W is issued which notifies the user that  
| the current level of command authorization is not being enforced.

| After message DSP1211W is issued, data sets without command authorization can  
| complete testing or debugging against the copy of the RECON data set. Message  
| DSP1211W also is issued after every subsequent command as a warning message.

---

## Tracking changes made to the RECON data sets

For administrative purposes, it is a good practice to keep a log of the changes made to the RECON data set. You can provide your own exit routine which is to be called every time RECON records are changed or read.

This exit routine, also called the RECON I/O exit routine (DSPCEXT0), allows you to keep track of changes to the RECON data set in the form of a journal.

**Related reading:** See *IMS Version 11 Exit Routines* for detailed information about the RECON I/O exit routine.

---

## Chapter 42. Hints and tips for DBRC

This topic provides task-oriented instructions for frequently-used procedures, including the following topics.

---

### Changing the RECON data sets to output time stamps in local time of origin

If the listing of the RECON data set does not list the offset value for time stamps, the record might have been created when the offset value was different from the current offset value. Using DBRC commands to change or delete such a record will result in a “record-not-found” condition because the current offset is applied if it is not included in the time stamp parameter of the command.

To change the RECON data set to output time stamps in local time of origin, when they were created, use one of the following:

```
CHANGE.RECON TIMEFMT(0,0,P,4,PERM)
```

or

```
CHANGE.RECON TIMEFMT(0,0,P,2,PERM)
```

This will permanently change the RECON data set to output time stamps when the records were originally created with the appropriate offsets in punctuated four-digit or two-digit year formats. The time stamp listed can be used in commands such as the following:

```
CHANGE.PRILOG STARTIME('2007.063 16:11:58.123456 -08:00') -  
DSSTART('2007.063 16:11:58.123456 -08:00') ERROR RLDS
```

---

### Locating the last SLDS stop time in the RECON data set

Use the following procedure to find the last SLDS stop time in the RECON data set using the GENJCL.USER command. This procedure can be used while IMS is still running and can also be used if the PRISLD is already closed.

1. Create a skeletal JCL execution member.

Here is an example of a member named SELSLDS:

```
%SELECT SLDS(%USSID, LAST)  
%ENDSEL  
SLDSDD BEG=%SLDSTIM  
END=%SLDETIM  
VOL=%SLDVOLS  
UNT=%SLDUNIT  
DSN=%SLSDSN
```

2. Issue the GENJCL.USER command indicating the execution member to run with.

Below is a sample of the GENJCL.USER command indicating that the command is to be run with member SELSLDS:

```
GENJCL.USER MEMBER(SELSDS) NOJOB USERKEYS((%USSID, 'IMS1')) -  
JCLOUT(SYSPRINT)
```

Below is the SLDS listing:

```

PRISLD                                RECORD SIZE=    1104
START = 07.064 16:58:06.000000      *  SSID=IMS1    VERSION=10.1
STOP  = 07.064 17:00:04.656844      #DSN=6
GSGNAME=**NULL**
FIRST RECORD ID= 0000000000000001  PRILOG TOKEN= 0

DSN=IMSVS.SLDSP.IMS1.D07064.T1658060.V00      UNIT=SYSDA
START = 07.064 16:58:06.000000      FIRST DS LSN= 0000000000000001
STOP  = 07.064 16:58:16.108103      LAST  DS LSN= 0000000000000128
FILE SEQ=0001      #VOLUMES=0001
CHECKPOINT TYPES=80: SIMPLE=Y SNAPQ=N DUMPQ=N PURGE=N FREEZE=N

VOLSER=000000 STOPTIME = 07.064 16:58:16.108103
CKPTCT=1      CHKPT ID = 00.000 00:00:00.000000
LOCK SEQUENCE#= 000000000000

DSN=IMSVS.SLDSP.IMS1.D07064.T1658161.V00      UNIT=SYSDA
START = 07.064 16:58:16.108103      FIRST DS LSN= 0000000000000129
STOP  = 07.064 16:59:00.886236      LAST  DS LSN= 0000000000000BE9
FILE SEQ=0001      #VOLUMES=0001
CHECKPOINT TYPES=00: SIMPLE=N SNAPQ=N DUMPQ=N PURGE=N FREEZE=N

VOLSER=000000 STOPTIME = 07.064 16:59:00.886236
CKPTCT=0      CHKPT ID = 07.064 16:58:06.603513
LOCK SEQUENCE#= 000000000000

DSN=IMSVS.SLDSP.IMS1.D07064.T1659008.V00      UNIT=SYSDA
START = 07.064 16:59:00.886236      FIRST DS LSN= 0000000000000BEA
STOP  = 07.064 16:59:15.330532      LAST  DS LSN= 000000000000111A
FILE SEQ=0001      #VOLUMES=0001
CHECKPOINT TYPES=00: SIMPLE=N SNAPQ=N DUMPQ=N PURGE=N FREEZE=N

VOLSER=000000 STOPTIME = 07.064 16:59:15.330532
CKPTCT=0      CHKPT ID = 07.064 16:58:06.603513
LOCK SEQUENCE#= 000000000000

DSN=IMSVS.SLDSP.IMS1.D07064.T1659153.V00      UNIT=SYSDA
START = 07.064 16:59:15.330532      FIRST DS LSN= 000000000000111B
STOP  = 07.064 16:59:28.219762      LAST  DS LSN= 00000000000016A5
FILE SEQ=0001      #VOLUMES=0001
CHECKPOINT TYPES=00: SIMPLE=N SNAPQ=N DUMPQ=N PURGE=N FREEZE=N

VOLSER=000000 STOPTIME = 07.064 16:59:28.219762
CKPTCT=0      CHKPT ID = 07.064 16:58:06.603513
LOCK SEQUENCE#= 000000000000

07.064 17:00:18.768831      LISTING OF RECON      PAGE 0005
DSN=IMSVS.SLDSP.IMS1.D07064.T1659282.V01      UNIT=SYSDA
START = 07.064 16:59:28.219762      FIRST DS LSN= 00000000000016A6
STOP  = 07.064 17:00:04.440567      LAST  DS LSN= 0000000000001907
FILE SEQ=0001      #VOLUMES=0001
CHECKPOINT TYPES=08: SIMPLE=N SNAPQ=N DUMPQ=N PURGE=N FREEZE=Y

VOLSER=000000 STOPTIME = 07.064 17:00:04.440567
CKPTCT=1      CHKPT ID = 07.064 16:58:06.603513
LOCK SEQUENCE#= 000000000000

DSN=IMSVS.SLDSP.IMS1.D07064.T1659282.V02      UNIT=SYSDA
START = 07.064 17:00:04.440567      FIRST DS LSN= 0000000000001908
STOP  = 07.064 17:00:04.656844      LAST  DS LSN= 00000000000019A0
FILE SEQ=0001      #VOLUMES=0001
CHECKPOINT TYPES=00: SIMPLE=N SNAPQ=N DUMPQ=N PURGE=N FREEZE=N

VOLSER=000000 STOPTIME = 07.064 17:00:04.656844
CKPTCT=0      CHKPT ID = 07.064 17:00:04.417809
LOCK SEQUENCE#= 000000000000

```



Below is the output that results from issuing the previous sample GENJCL.USER command:

```
SLDSDD  BEG=07064170004440567-0800
        END=07064170004656844-0800
        VOL=000000
        UNT=SYSDA
        DSN=IMSVS.SLDSP. IMS1.D07064.T1659282.V02
```

---

## Adjusting GENMAX when it is reached or it is too high

Prior to image copy utility execution, verification exit processing may determine that one of the following situations is true when image copy data sets are being reused. In such cases, the GENMAX value is either too high or too low. In either case, a command must be issued to correct the problem.

- **Situation 1-** GENMAX value is too low
  - The GENMAX value was reached.
  - The oldest image copy cannot be reused because it is within the recovery period.
  - No available image copies exist that can be used.
  - Message DSP0063I was displayed.
  - Increase GENMAX and define additional image copy data sets.
  - Change the recovery period to allow the oldest image copy data set to be used.

Issue the CHANGE.DBDS command to increase the GENMAX and INIT.IC parameters to define the additional image copy data sets:

1. Create a JCL job similar to one of the following jobs. These examples assume that your DBD name is THISDB, that your area name is AREA1, and that your previous GENMAX value was 2.

```
//CHNGDBDS JOB
//SYSIN DD *
    CHANGE.DBDS DBD(THISDB) AREA(AREA1) -
                GENMAX(4)
/*//INITIC JOB
//SYSIN DD *
    INIT.IC DBD(THISDB) DDN(DDN1) -
            ICDSN(IMS.*.NEWICDSN)
    INIT.IC DBD(THISDB) DDN(DDN1) -
            ICDSN(IMS.*.NEWICDSN2)
/*
```

Issue CHANGE.DBDS RECOVPD to reduce the recovery period.

```
//CHNGDBDS JOB
//SYSIN DD *
    CHANGE.DBDS DBD(THISDB) AREA(AREA1) RECOVPD(10)
/*
```

2. Run the job.
3. Check the JES log to ensure that the job ran successfully.
4. Run a LIST command to see the new GENMAX value (optional).
5. Run your image copy job.

- **Situation 2 -** GENMAX value is too high
  - GENMAX has not been reached so the in use image copy data sets cannot be used.
  - No available image copy data sets exist that can be used.



- Message DSP0084I was displayed.
- No recovery period is defined.
- Issue an INIT.IC command to define a new image copy data set for the identified database data set or area data set

**-OR-**

- Issue a CHANGE.DBDS command to lower the GENMAX value.

**Attention:** If GENMAX is lowered using the CHANGE.DBDS command, the new GENMAX value is recorded regardless of whether the oldest image copies are able to be deleted or not, because they are within the recovery period.

#### Related reference:

 CHANGE.DBDS command (Commands)

## Closing an open online PRILOG record

During time stamp recoveries, when DBRC reports that there are open allocations that are not needed for the recovery, it might be easier to close the open PRILOG record rather than deleting the individual allocations. Use the NOTIFY.PRILOG command to close the open PRILOG record.

If an open PRILOG record is found for an online IMS subsystem, and the PRILOG record is not for the current run of IMS, it indicates that the last OLDS for this online IMS has not yet been archived. If the OLDS is no longer available and you must close the open PRILOG record, the following commands can be used to create a dummy DSN entry in the PRILOG:

```
/* This command creates dummy dataset */
NOTIFY.PRILOG DSN(DUMMY) STARTIME('2010.072 10:23:53.979935 -07:00') -
    FIRSTREC(00) RLDS SSID(IMS1) VOLSER(VOL001)

/* This command closes PRILOG record-Stoptime of 0's becomes runtime */
NOTIFY.PRILOG RUNTIME('2010.072 11:23:53.980001 -07:00') -
    STARTIME('2010.072 10:23:53.979935 -07:00') -
    LASTREC(500) RLDS SSID(IMS1)

/* This command changes last PRILOG, DUMMY, to ERROR */
CHANGE.PRILOG STARTIME('2010.072 10:23:53.979935 -07:00') ERROR -
    DSSTART ('2010.072 10:24:43.286431 -07:00')
```

The following example is the PRILOG record as it appears in the RECON data set before issuing the NOTIFY.PRILOG command to close it:

```
PRILOG                                RECORD SIZE=      464
START = 2010.072 10:23:53.979935 -07:00*  SSID=IMS1  VERSION=11.1
STOP  = 0000.000 00:00:00.000000 -07:00    #DSN=2
GSGNAME=**NULL**
FIRST RECORD ID= 0000000000000001    PRILOG TOKEN= 0
EARLIEST CHECKPOINT = 2010.072 10:23:54.517056 -07:00

DSN=IMSVS.RLDSP.IMS1.D07072.T1023539.V00                                UNIT=SYSDA
START = 2010.072 10:23:53.979935 -07:00  FIRST DS LSN= 0000000000000001
STOP  = 2010.072 10:24:32.697834 -07:00  LAST  DS LSN= 00000000000006D1
FILE SEQ=0001    #VOLUMES=0001

VOLSER=000000  STOPTIME = 2010.072 10:24:32.697834 -07:00
CKPTCT=1    CHKPT ID = 2010.072 10:23:54.517056 -07:00
LOCK SEQUENCE#= 000000000000

DSN=IMSVS.RLDSP.IMS1.D07072.T1024326.V00                                UNIT=SYSDA
START = 2010.072 10:24:32.697834 -07:00  FIRST DS LSN= 00000000000006D2
STOP  = 2010.072 10:24:43.286431 -07:00  LAST  DS LSN= 0000000000000831
```

```

FILE SEQ=0001      #VOLUMES=0001

VOLSER=000000 STOPTIME = 2010.072 10:24:43.286431 -07:00
CKPTCT=0      CHKPT ID = 2010.072 10:23:54.517056 -07:00
LOCK SEQUENCE#= 000000000000

LOGALL
START      = 2010.072 10:23:53.979935 -07:00*
EARLIEST ALLOC TIME = 2010.072 10:24:23.704405 -07:00
DBDS ALLOC=1                      -DBD-      -DDN-      -ALLOC-
                                      DBVHDJ05 CJVHDG1E 1

```

The following example is the PRILOG record as it appears in the RECON data set after issuing the NOTIFY.PRILOG command to close it (differences are highlighted).

```

PRILOG                                RECORD SIZE=      624
START = 2010.072 10:23:53.979935 -07:00*  SSID=IMS1      VERSION=11.1
STOP  = 2010.072 11:23:53.980001 -07:00    #DSN=3
GSGNAME=**NULL**
FIRST RECORD ID= 0000000000000001      PRILOG TOKEN= 0
EARLIEST CHECKPOINT = 2010.072 10:23:54.517056 -07:00

DSN=IMSVS.RLDSP.IMS1.D07072.T1023539.V00      UNIT=SYSDA
START = 2010.072 10:23:53.979935 -07:00 FIRST DS LSN= 0000000000000001
STOP  = 2010.072 10:24:32.697834 -07:00 LAST  DS LSN= 00000000000006D1
FILE SEQ=0001      #VOLUMES=0001

VOLSER=000000 STOPTIME = 2010.072 10:24:32.697834 -07:00
CKPTCT=1      CHKPT ID = 2010.072 10:23:54.517056 -07:00
LOCK SEQUENCE#= 000000000000

DSN=IMSVS.RLDSP.IMS1.D07072.T1024326.V00      UNIT=SYSDA
START = 2010.072 10:24:32.697834 -07:00 FIRST DS LSN= 00000000000006D2
STOP  = 2010.072 10:24:43.286431 -07:00 LAST  DS LSN= 0000000000000831
FILE SEQ=0001      #VOLUMES=0001

VOLSER=000000 STOPTIME = 2010.072 10:24:43.286431 -07:00
CKPTCT=0      CHKPT ID = 2010.072 10:23:54.517056 -07:00
LOCK SEQUENCE#= 000000000000

DSN=DUMMY                                ERROR UNIT=3400
START = 2010.072 10:24:43.286431 -07:00 FIRST DS LSN= 0000000000000000
STOP  = 2010.072 11:23:53.980001 -07:00 LAST  DS LSN= 0000000000001F4
FILE SEQ=0001      #VOLUMES=0001

VOLSER=VOL001 STOPTIME = 2010.072 11:23:53.980001 -07:00
CKPTCT=0      CHKPT ID = 0000.000 00:00:00.000000 -07:00
LOCK SEQUENCE#= 30E100000000

LOGALL
START      = 2010.072 10:23:53.979935 -07:00*
EARLIEST ALLOC TIME = 2007.072 10:24:23.704405 -07:00
DBDS ALLOC=1                      -DBD-      -DDN-      -ALLOC-
                                      DBVHDJ05 CJVHDG1E 1

```

#### Related reference:

-  NOTIFY.PRILOG command (for RLDS) (Commands)
-  LIST.LOG command (for a PRILOG family) (Commands)

---

## Working with subsystem records (SSYS)

If authorization for a database fails and you list the database and all the subsystem records in the RECON data set, you might wonder how the subsystem names (SSIDs) were created.

The following list discusses different aspects of the SSIDs.

- Naming conventions for SSIDs in RECON subsystem records

The SSYS record is created when a signon call is issued to DBRC. The SSID must be unique. Normal conventions for creating the subsystem name are:

IMS and DBCTL subsystems = IMSID value from IMSCTRL SYSGEN macro  
(four characters)

BATCH and UTILITY SUBSYSTEMS = JOBNAME

XRF = RSENAME (Recovery Service Element) of the IMS systems  
(active and alternate)

Exceptions include the following:

ONLINE IMAGE COPY - XXXTZZZZ where

XXXT is the DMB number of the database + DCB number translated  
to 0-9,A-Z  
ZZZZ is the control region IMSID

Online image copy is the only BMP that creates its own subsystem record in the RECON data set.

- Batch backout

If batch backout is processing a log created by a batch subsystem, the SSID that is used is the job name of the subsystem that is being backed out. The job name of the batch backout utility job is not used.

If batch backout is processing a log that was created by an online subsystem, the job name of the backout utility job is used as the SSID.

- Deleting a subsystem record

When you no longer need a particular subsystem record in the RECON data set, the following commands can be used to delete the SSYS:

```
CHANGE.SUBSYS SSID(XXXXX) STARTRCV
CHANGE.SUBSYS SSID(XXXXX) ENDRECOV <---(removes auth from DB)
DELETE.SUBSYS SSID(XXXXX)
```

- Subsystem record size

The subsystem record grows in size. If the subsystem record grows to exceed the 16M maximum record size, IMS abends.

---

## Removing authorization inconsistency between the SSYS from DB/AREA records

To change database authorization or if authorization fails, a LIST.DB command shows that the DB or area is authorized to a nonexistent subsystem and backout is not required.

Use the following command to remove the authorization. You can also use the command below to remove a database/area name from the list in a subsystem record when a LIST.SUBSYS shows it authorized and a LIST.DB does not show it authorized.

```
CHANGE.DB DBD(dbname) SSID(XXXX) UNAUTH (or)
CHANGE.DB DBD(dbname) AREA(areaname) SSID(XXXX) UNAUTH
```

---

## Restarting the change accumulation process for logs

Change accumulation erroneously stops processing logs and the source of the problem is traced to an OLDS left in the ARC STARTED state. An IPL of z/OS is done on the CPU where the archive job was run and the status of the OLDS remains in this state.

GENJCL.ARCHIVE jobs do not include any OLDS in ARC STARTED status. The online IMS that created the OLDS is not running on this CPU.

If the online IMS is run on the same CPU as the archive job when IMS is restarted, it resets the status of all OLDS in ARC STARTED status to ARC NEEDED. The OLDS is included in the next archive job. To avoid this problem, run archive jobs on the same CPU as their corresponding IMS systems.

In cases where manual intervention is required, use the following command to change the status of the OLDS to archive needed (ARNEEDED):

```
CHANGE.PRILOG OLDS(DFSOLPXX) ARNEEDED SSID(XXXX)
```

---

## Restarting the change accumulation work when it states nothing to process

Change accumulation has issued a message stating that it found no logs to process. A review of the listing reveals that there are many logs that need to be processed.

DBRC selects all logs needed to satisfy the allocations for each DBDS since the effective image copy time. The logs are put in log volume start time order as a way for DBRC to keep track of which logs it has processed. DBRC processes all available logs and truncates the list of log volumes when it encounters a log in error, an open log, or when it detects an unarchived OLDS. A message indicating the condition that was encountered (such as an open log) is issued.

If the last change accumulation execution record shows that a SUBSET of logs was input to it last, the STOPTIME reflects the start time of the first unavailable log volume. Find this log volume time stamp in the listing of the RECON data set. If no log volume exists with this start time, there is an unarchived OLDS.

The following timeline illustrates this open OLDS condition:

```
T1--T2---T3--T4-T5-----T6-----T7-----T8----T9---  
  
|---A1---D1--|-----A4-----// (open)  
Log1 (DSN1)   (OPEN OLDS)  
  
      |--A2--|  |--A3-|  
      Log2   Log3
```

If GENJCL.CA is run at T7, it includes Log1(DSN1) and has a STOPTIME of T4 and is a SUBSET. T4 is the start time of the first unavailable log. After this JCL is executed, any attempts to run GENJCL.CA result in issuing two messages; one stating that an unarchived OLDS exists and another stating that there are no logs to process until the open OLDS is closed and archived.

---

## Moving log data sets

DBRC records information about IMS logs in the RECON data set. This information includes the log data set name (DSN), its starting and ending times, and the volume serial (VOLSER).

The best way to avoid lost logs is to use cataloged log data sets and the DBRC CATDS option.

For non-cataloged log data sets, inform DBRC about changes by using the CHANGE.PRILOG or CHANGE.SECLOG command.

**Related reading:** See the CHANGE.PRILOG and the CHANGE.SECLOG commands in *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* .

---

## Cataloging data sets

You can indicate whether image copy, change accumulation, and log data sets are cataloged or SMS managed by using the CATDS or NOCATDS keywords on the CHANGE.RECON or INIT.RECON command.

The CHANGE.RECON and INIT.RECON commands update the RECON data set header record accordingly if you specify either CATDS or NOCATDS. The only difference between the two commands in respect to cataloging is that INIT.RECON can be issued only once for each set of RECON data sets (3) required for an IMS region.

To specify that these data sets are cataloged specify:

```
//CHGRECON JOB
:
//SYSIN DD *
        CHANGE.RECON CATDS
/*
```

**-OR-**

```
//INITRCON JOB
:
//SYSIN DD *
        INIT.RECON CATDS
/*
```

The data set must have been initialized as cataloged for CATDS to be effective with the CHANGE.RECON command.

Specifying NOCATDS on the CHANGE.RECON or INIT.RECON command establishes that these data sets, regardless of their cataloged status, are not to be treated as cataloged.

To specify that these data sets are not to be treated as cataloged specify:

```
//CHGRECON JOB
:
//SYSIN DD *
      CHANGE.RECON NOCATDS
/*
```

**-OR-**

```
//INITRCON JOB
:
//SYSIN DD *
      INIT.RECON NOCATDS
/*
```

**Related reading:** See *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands* for more information.

---

## Performing multiple cold starts in a test environment

Many times in a test environment, you may want to cold start IMS. In order to cold start IMS the last OLDS must be closed.

If you have no need to close the OLDS, you can use the following commands to close the OLDS and mark it archived in the RECON data set so it can be reused:

```
NOTIFY.PRILOG STARTIME(%OLDOTIM) RUNTIME(070682226022-0800) -
      LASTREC(670) OLDS(DFSOLP02) SSID(IMS1)
CHANGE.PRILOG OLDS(DFSOLP02) ARCHIVED SSID(IMS1)
```

The entries for each OLDS (such as: DFSOLP00, DSPOLP01, and DSPOLD02) in the PRIOLD record are built when the OLDSs are used (if they do not already exist in the RECON data set). If you also need to delete the OLDS from the RECON data set, the following commands can be used:

```
DELETE.LOG OLDS(DFSOLP00) SSID(IMS1)
DELETE.LOG OLDS(DFSOLP01) SSID(IMS1)
DELETE.LOG OLDS(DFSOLP02) SSID(IMS1) LASTCLOS
```

Note that the LASTCLOS is necessary to delete the last OLDS used by IMS. The PRIOLD record is also deleted when the last DD name entry is removed.

The PRIOLD record before issuing the commands:

```
PRIOLD
  SSID=IMS1          # DD ENTRIES=3
  EARLIEST CHECKPOINT = 07.068 20:55:58.622862

  DDNAME=DFSOLP00   DSN=IMSTESTL.IMS01.OLDSP0
  START = 07.068 20:55:57.984378      FIRST DS LSN= 0000000000000001
  STOP  = 07.068 20:56:39.570062      LAST  DS LSN= 000000000000006D1
  LOCK SEQUENCE# = 000000000000
  STATUS=ARC COMPLT                                FEOV=YES   AVAIL
  PRILOG TIME=07.068 20:55:57.984378      ARCHIVE JOB NAME=JT205639
  VERSION=10.1

  DDNAME=DFSOLP01   DSN=IMSTESTL.IMS01.OLDSP1
  START = 07.068 20:56:39.570062      FIRST DS LSN= 000000000000006D2
  STOP  = 07.068 20:56:53.689006      LAST  DS LSN= 00000000000000831
  LOCK SEQUENCE# = 000000000000
  STATUS=ARC COMPLT                                FEOV=YES   AVAIL
  PRILOG TIME=07.068 20:55:57.984378      ARCHIVE JOB NAME=JT205653
  VERSION=10.1
```

```
DDNAME=DFSOLP02 DSN=IMSTESTL.IMS01.OLDSP2
START = 07.068 20:56:53.689006 FIRST DS LSN= 0000000000000832
STOP = 00.000 00:00:00.000000 LAST DS LSN= 0000000000000000
LOCK SEQUENCE# = 000000000000
STATUS=ACTIVE FEOV=NO AVAIL
PRILOG TIME=07.068 20:55:57.984378
VERSION=10.1
```

The PRIOLD record after issuing the commands to close and mark DFSOLP02 as archived (with the differences highlighted):

```
PRIOLD
SSID=IMS1 # DD ENTRIES=3
EARLIEST CHECKPOINT = 07.068 20:55:58.622862

DDNAME=DFSOLP00 DSN=IMSTESTL.IMS01.OLDSP0
START = 07.068 20:55:57.984378 FIRST DS LSN= 0000000000000001
STOP = 07.068 20:56:39.570062 LAST DS LSN= 00000000000006D1
LOCK SEQUENCE# = 000000000000
STATUS=ARC COMPLT FEOV=YES AVAIL
PRILOG TIME=07.068 20:55:57.984378 ARCHIVE JOB NAME=JT205639
VERSION=10.1

DDNAME=DFSOLP01 DSN=IMSTESTL.IMS01.OLDSP1
START = 07.068 20:56:39.570062 FIRST DS LSN= 00000000000006D2
STOP = 07.068 20:56:53.689006 LAST DS LSN= 0000000000000831
LOCK SEQUENCE# = 000000000000
STATUS=ARC COMPLT FEOV=YES AVAIL
PRILOG TIME=07.068 20:55:57.984378 ARCHIVE JOB NAME=JT205653
VERSION=10.1

DDNAME=DFSOLP02 DSN=IMSTESTL.IMS01.OLDSP2
START = 07.068 20:56:53.689006 FIRST DS LSN= 0000000000000832
STOP = 07.068 22:26:02.200000 LAST DS LSN= 000000000000029E
LOCK SEQUENCE# = 000000000000
STATUS=ARC COMPLT FEOV=NO AVAIL
PRILOG TIME=07.068 20:55:57.984378
VERSION=10.1
```

---

## Reducing potential RECON data set enqueue problems

The following list describes some scenarios that can cause the RECON data set enqueue problems and contain suggestions for avoiding these problems.

- Shared DASD environment scenarios

Operating in a shared DASD environment, the most common cause of RECON data set enqueue problems is failing to follow the recommendation to catalog each RECON data set in its own ICF catalog on the same volume as the RECON data set.

If you use serial access and if you catalog each RECON data set in its own ICF catalog on the same volume as the RECON data set and still have problems; examine your GRS, (or equivalent) RESERVE conversion list, to determine how you process SYSIGGV2 and DSPURI01 QNAMEs. A couple of combinations might lead to deadlocks.

- Non-shared DASD environment scenarios

If you are operating in a non-shared DASD environment and are having problems, this is probably not caused by deadlock but rather by contention and slow performance. Here are a few things to look at in this situation:

- Minimize the amount of time any job holds the RECON data set.

One way to minimize that time is to tune the LSR buffer pool DBRC uses when accessing the RECON data sets. There is a CSECT, DSPBUFFS (serial

access only), that contains the values used for online and batch. See “Buffer Size Specification facility (DSPBUFFS)” in *IMS Version 11 Exit Routines* for information about how to zap it and change the values. The defaults might be low for your usage. For parallel access, the IDGSMsxx RLS\_MAX\_POOL\_SIZE definitions should be looked into instead.

- Analyze the other contents of the volumes where the RECON data set resides. Consider isolating these volumes, to prevent interference from other I/O activity (and vice versa). Consider placing the RECON data sets on high performance cached DASD, or perhaps solid state DASD.

**Related reading:** See “Avoiding RECON data set deadlock situations for serial access mode” on page 538 for more information.





---

## Chapter 43. Resource names for command authorization

DBRC uses a list of resource names for command authorization support.

The following table lists the resources that can be protected.

*Table 45. Resource names for command authorization*

| Verb   | Modifier       | Qualifier |
|--------|----------------|-----------|
| AUTH   | not applicable | .dbname   |
| BACKUP | .RECON         |           |
| CHANGE | .ADS           | .dbname   |
| CHANGE | .BKOUT         | .ssid     |
| CHANGE | .CA            | .grpname  |
| CHANGE | .CAGRP         | .grpname  |
| CHANGE | .DB            | .dbname   |
| CHANGE | .DB            | .ALL      |
| CHANGE | .DBDS          | .dbname   |
| CHANGE | .DBDSGRP       | .grpname  |
| CHANGE | .IC            | .dbname   |
| CHANGE | .PART          | .dbname   |
| CHANGE | .PRILOG        | .OLDS     |
| CHANGE | .PRILOG        | .RLDS     |
| CHANGE | .PRILOG        | .SLDS     |
| CHANGE | .PRILOG        | .TSLDS    |
| CHANGE | .RECON         |           |
| CHANGE | .RECON         | .CATDS    |
| CHANGE | .RECON         | .NOCATDS  |
| CHANGE | .RECON         | .CHECK17  |
| CHANGE | .RECON         | .CHECK44  |
| CHANGE | .RECON         | .NOCHECK  |
| CHANGE | .RECON         | .CMDAUTH  |
| CHANGE | .RECON         | .DASDUNIT |
| CHANGE | .RECON         | .DUAL     |
| CHANGE | .RECON         | .REPLACE  |
| CHANGE | .RECON         | .FORCER   |
| CHANGE | .RECON         | .NOFORCER |
| CHANGE | .RECON         | .IMSPLEX  |
| CHANGE | .RECON         | .NOPLEX   |
| CHANGE | .RECON         | .LISTDL   |
| CHANGE | .RECON         | .NOLISTDL |
| CHANGE | .RECON         | .LOGALERT |

Table 45. Resource names for command authorization (continued)

| Verb    | Modifier | Qualifier  |
|---------|----------|------------|
| CHANGE  | .RECON   | .LOGRET    |
| CHANGE  | .RECON   | .MINVERS   |
| CHANGE  | .RECON   | .SIZALERT  |
| CHANGE  | .RECON   | .SSID      |
| CHANGE  | .RECON   | .STARTNEW  |
| CHANGE  | .RECON   | .NONEW     |
| CHANGE  | .RECON   | .TAPEUNIT  |
| CHANGE  | .RECON   | .TIMEZONE  |
| CHANGE  | .RECON   | .TIMEZIN   |
| CHANGE  | .RECON   | .TIMEFMT   |
| CHANGE  | .RECON   | .TRACEON   |
| CHANGE  | .RECON   | .TRACEOFF  |
| CHANGE  | .RECON   | .UPGRADE   |
| CHANGE  | .SG      | .gsgname   |
| CHANGE  | .SUBSYS  | .ssid      |
| CHANGE  | .SUBSYS  | .ALL       |
| CHANGE  | .UIC     | .dbname    |
| CLEANUP | .RECON   |            |
| CLEANUP | .RECON   | .RETPRD    |
| CLEANUP | .RECON   | .TIME      |
| CLEANUP | .RECON   | .DBRANGE   |
| CLEANUP | .RECON   | .DBONLY    |
| CLEANUP | .RECON   | .LASTIC    |
| CLEANUP | .RECON   | .LISTDL    |
| CLEANUP | .RECON   | .NOLISTDL  |
| DELETE  | .ADS     | .dbname    |
| DELETE  | .ALLOC   | .dbname    |
| DELETE  | .BKOUT   | .ssid      |
| DELETE  | .CA      | .grpname   |
| DELETE  | .CAGRP   | .grpname   |
| DELETE  | .DB      | .dbname    |
| DELETE  | .DBDS    | .dbname    |
| DELETE  | .DBDSGRP | .grpname   |
| DELETE  | .GSG     | .gsgname   |
| DELETE  | .IC      | .dbname    |
| DELETE  | .LOG     | .INACTIVE  |
| DELETE  | .LOG     | .OLDS      |
| DELETE  | .LOG     | .STARTTIME |
| DELETE  | .LOG     | .TOTIME    |
| DELETE  | .PART    | .dbname    |

Table 45. Resource names for command authorization (continued)

| Verb   | Modifier | Qualifier |
|--------|----------|-----------|
| DELETE | .RECOV   | .dbname   |
| DELETE | .REORG   | .dbname   |
| DELETE | .SG      | .gsgname  |
| DELETE | .SUBSYS  | .ssid     |
| DELETE | .UIC     | .dbname   |
| GENJCL | .ARCHIVE | .ssid     |
| GENJCL | .CA      | .grpname  |
| GENJCL | .CLOSE   | .ssid     |
| GENJCL | .IC      | .dbname   |
| GENJCL | .IC      | .grpname  |
| GENJCL | .OIC     | .dbname   |
| GENJCL | .OIC     | .grpname  |
| GENJCL | .RECEIVE | .dbname   |
| GENJCL | .RECEIVE | .grpname  |
| GENJCL | .RECOV   | .dbname   |
| GENJCL | .RECOV   | .grpname  |
| GENJCL | .USER    | .mbrname  |
| GENJCL | .USER    | .dbname   |
| GENJCL | .USER    | .grpname  |
| INIT   | .ADS     | .dbname   |
| INIT   | .CA      | .grpname  |
| INIT   | .CAGRP   | .grpname  |
| INIT   | .DB      | .dbname   |
| INIT   | .DBDS    | .dbname   |
| INIT   | .DBDSGRP | .grpname  |
| INIT   | .GSG     | .gsgname  |
| INIT   | .IC      | .dbname   |
| INIT   | .PART    | .dbname   |
| INIT   | .RECON   |           |
| INIT   | .SG      | .gsgname  |
| LIST   | .BKOUT   | .ssid     |
| LIST   | .BKOUT   | .ALL      |
| LIST   | .CAGRP   | .grpname  |
| LIST   | .CAGRP   | .ALL      |
| LIST   | .DB      | .dbname   |
| LIST   | .DB      | .ALL      |
| LIST   | .DB      | .TYPEIMS  |
| LIST   | .DB      | .TYPEFP   |
| LIST   | .DB      | .TYPHALDB |
| LIST   | .DBDS    | .dbname   |

Table 45. Resource names for command authorization (continued)

| Verb   | Modifier       | Qualifier |
|--------|----------------|-----------|
| LIST   | .DBDS          | .grpname  |
| LIST   | .DBDSGRP       | .grpname  |
| LIST   | .DBDSGRP       | .ALL      |
| LIST   | .GSG           | .gsgname  |
| LIST   | .GSG           | .ALL      |
| LIST   | .HISTORY       | .dbname   |
| LIST   | .HISTORY       | .grpname  |
| LIST   | .LOG           | .ALL      |
| LIST   | .LOG           | .ALLOLDS  |
| LIST   | .LOG           | .OLDS     |
| LIST   | .LOG           | .STARTIME |
| LIST   | .RECON         |           |
| LIST   | .RECON         | .STATUS   |
| LIST   | .SUBSYS        | .ONLINE   |
| LIST   | .SUBSYS        | .ssid     |
| LIST   | .SUBSYS        | .ALL      |
| LIST   | .SUBSYS        | .BATCH    |
| NOTIFY | .ALLOC         | .dbname   |
| NOTIFY | .BKOUT         | .ssid     |
| NOTIFY | .CA            | .grpname  |
| NOTIFY | .IC            | .dbname   |
| NOTIFY | .PRILOG        | .OLDS     |
| NOTIFY | .PRILOG        | .RLDS     |
| NOTIFY | .PRILOG        | .SLDS     |
| NOTIFY | .PRILOG        | .TSLDS    |
| NOTIFY | .PRILOG        | .DPROP    |
| NOTIFY | .RECOV         | .dbname   |
| NOTIFY | .REORG         | .dbname   |
| NOTIFY | .SECLOG        | .OLDS     |
| NOTIFY | .SECLOG        | .RLDS     |
| NOTIFY | .SECLOG        | .SLDS     |
| NOTIFY | .SECLOG        | .TSLDS    |
| NOTIFY | .SECLOG        | .DPROP    |
| NOTIFY | .SUBSYS        | .ssid     |
| NOTIFY | .UIC           | .dbname   |
| RESET  | .GSG           | .gsgname  |
| STDBRC | not applicable | .ssid     |

For more information about command authorization, see “DBRC security” on page 563.

---

## Part 4. IMS system recovery

| An Extended Recovery Facility (XRF), a Remote Site Recovery (RSR) complex, or  
| an IMSplex allows you to maintain continuity of online transaction processing by  
| giving you the ability to rapidly resume end-user service when a failure occurs.



---

## Chapter 44. Extended Recovery Facility Overview

An Extended Recovery Facility (XRF) complex allows you to maintain continuity of online transaction processing by giving you the ability to rapidly resume end-user service when a failure occurs.

### Related concepts:

“Running the Extended Recovery Facility” on page 10

“Recovery in an XRF complex” on page 114

 IMS recovery using Extended Recovery Facility (Operations and Automation)

---

### Overview of XRF

The XRF is a set of functions within existing programs that can be used by an installation to achieve a high level of availability, reliability, and continuity of service for selected end users. XRF responds to the need for continuity of online transaction processing.

The required set of IBM software products that work together with IMS to make up XRF are:

- z/OS
- DFSMS
- VTAM
- Network Control Program (NCP) for an XRF complex that uses USERVAR
- System Support Programs (SSP) for an XRF complex that uses USERVAR

The XRF approach to high availability is based on two assumptions:

- Many IMS installations try to minimize both planned and unplanned outages. These installations are willing to devote extra resources to improve the service for their high-priority work.
- A defect that causes a failure in one environment does not necessarily cause a failure in a different environment.

XRF does not eliminate outages. Even if all unplanned outages (such as hardware and software failures) disappeared, planned outages (such as maintenance, configuration changes, and migration) still occur. Although XRF does not provide continuous availability, it does reduce the impact of planned and unplanned outages on end users.

The remote site recovery (RSR) feature provides many of the same benefits that XRF provides. For a comparison of XRF and RSR, see Table 55 on page 680.

### XRF in an IMSplex

XRF can be used in an IMSplex in the same way that it is used on a local system. The XRF alternate must be defined in the IMSplex and have access to the Common Service Layer (CSL).

A Structured Call Interface (SCI) must reside on the z/OS image where the XRF alternate resides.



An XRF system registers for IMS commands with the Operations Manager (OM) using its IMSID. When commands are routed to an XRF active or alternate system, the IMSID should be used on the ROUTE().

For owned resources, an XRF system defines its resources to the Resource Manager (RM) using the RSENAME as the resource owner. When the XRF alternate system takes over, it takes over ownership of the resources.

You cannot use XRF to recover an entire IMSplex, but you can use XRF to recover an individual IMS in an IMSplex.

HALDB OLR cannot be initiated on an XRF alternate system. For more information about HALDB OLR support, see *IMS Version 11 Database Administration*.

---

## Installation types that benefit from XRF

Among the data processing installations that are especially sensitive to disruptions in service are banking institutions, credit verification companies, and manufacturing-process control centers. In these and other installations, loss of service can cause end-user dissatisfaction and loss of business.

Because of the growth of interactive workload and critical business applications, availability has become a major concern for many IMS installations. An installation that uses IMS to log employees on and off their jobs during shift change, for example, needs high availability during two brief periods of each work day. A bank that uses IMS to process transactions needs high availability from 6 a.m. to 10 p.m. An installation that requires uninterrupted processing 24 hours a day, 7 days a week, regards outages of 1 to 2 hours as unacceptable. Although XRF cannot eliminate outages at these installations, it can remove much of the disruption to end users.

Generally, an installation that can benefit from XRF has:

- Multiple large-scale systems in a single operational environment
- Many end users served by IMS
- A requirement for very high availability to end users

Such an installation typically operates in a well-structured and tightly controlled environment, one where the effective installation management needed to implement XRF exists.

---

## XRF concepts and terminology

High availability with only short interruptions in end-user service can be difficult to achieve for IMS failures. Before restart can take place, termination processing might need to be performed and the cause of failure determined. Restart itself can be time consuming if regions must be restarted and terminal sessions reestablished.

An XRF complex allows you to rapidly resume end-user service when a failure occurs. The configuration consists of a primary subsystem (called the *active IMS*), usually operating in a preferred processor, and an alternate subsystem (called the *alternate IMS*), tracking the activities of the active IMS concurrently in the same or in a different processor.

The active IMS processes the IMS workload in the same way that it would without XRF. It does not do any more work than its normal processing and logging.

Through the log, it informs the alternate IMS of its activity. A combination of surveillance mechanisms using the log, the restart data set (RDS), and an ISC link, alert the alternate IMS to problems in the active IMS. The active IMS is not aware of the activities of the alternate IMS.

The alternate IMS does not process any of the IMS transactions that are entered in the active IMS; it monitors the active IMS for an indication of trouble. The alternate IMS records resource changes in the active IMS and updates its control blocks and buffers, continuously changing its status to match that of the active IMS. This alternate IMS is in a state of readiness so that work can quickly shift from the active IMS to the alternate IMS without waiting for dependent regions to be started, data sets to be allocated and opened, and sessions to be established. This shift of the workload from the active IMS to the alternate IMS is called a *takeover*. After a takeover has occurred, you can establish a new alternate IMS.

When an IMS failure occurs, or when failures that affect the z/OS operating system or entire central processor complex (CPC) occur, the alternate IMS assumes the workload of the active IMS. If end users are using appropriate terminals, they experience little disruption when a takeover occurs. The effect is of a single-system image; the end user does not need to know which IMS is being used and might not even know that a switch in the IMS host system occurred.

A takeover can occur when the active IMS abends. It can also occur for some of the other common causes of outages at an IMS installation, such as:

- Surveillance-detectable IMS failures
- Surveillance-detectable z/OS failures, loops, or wait states
- CPC failures
- VTAM failures that results in a TPEND exit
- Internal Resource Lock Manager (IRLM) failures that results in a STATUS exit

You can also invoke XRF to introduce planned changes into a production environment with minimal disruption to end users.

One of the major service elements is the *central processor complex (CPC)*. The CPC is a physical collection of hardware that consists of main storage, one or more central processors, timers, and channels. A CPC runs under the control of a single operating system. It can be either a uniprocessor or a multiprocessor (including a dyadic processor).

When discussing a 3705, 3725, or 3745 Communication Controller, which is required when your XRF complex uses USERVAR, this information uses the term *37x5 Communication Controller*.

**Note:** Installations that want to migrate off of the 3745 controller have the following options:

- Use the Communication Controller for Linux on System z (CCL) as a replacement for the 3745 controller to allowed continued use of XRF
- Migrate from XRF to VTAM Generic Resources (VGR), which requires a parallel sysplex environment
- Use IMS Multi-Node Persistent Sessions (MNPS) support for XRF. Due to performance limitations, this is not a recommended solution.

**Recommendation:** Use either the CCL or VGR as a permanent replacement solution for the 3745 controller.

The following terms are specific to XRF:

**recoverable service element (RSE)**

The active IMS system and the alternate IMS system that work as a unit in an XRF complex.

The two IMS systems in an RSE share an RSENAME and either an MNPS ACB name or USERVAR tables, depending on which method you choose to manage terminal sessions for XRF. You can start either IMS system in an RSE as the active IMS system (subject to operational restrictions), and then start the other as the alternate IMS system. Both identify themselves to the availability manager (AVM) component of z/OS .

**dependent service element (DSE)**

An element of the active IMS system that has a counterpart in the alternate IMS system but cannot trigger a takeover on its own.

A DSE depends on IMS to recognize when it fails and to request a takeover on its behalf. The CPC, z/OS, VTAM, and IRLM are DSEs.

**takeover**

A shift of workload from the active IMS system to the alternate IMS system.

Two kinds of takeovers are:

**planned takeover**

An intentional shift of the IMS workload to the alternate IMS system to allow maintenance of the active IMS system.

**unplanned takeover**

A shift of the IMS workload to the alternate IMS due to a failure of the active IMS system or the CPC, z/OS operating system, or VTAM associated with the active IMS system.

**DASD service element**

The DASD associated with an XRF complex.

The DASD service element is neither recoverable nor dependent; it is shared by both the active IMS system and the alternate IMS system. A major failure in the DASD service element might terminate service to end users. This is also true of the 37x5 Communication Controller, if your XRF complex uses USERVAR.

## XRF complexes

To use XRF in your installation, you must build a suitable configuration, called an XRF complex. An XRF complex consists of hardware (such as the CPCs and DASDs) and the licensed programs.

The following lists the required hardware and licensed programs required for XRF:

- Two CPCs capable of operating independently, such as two IBM eServer™ zSeries 990 CPCs
- z/OS (including DFSMS), IMS, and VTAM running in each CPC
- IMS system logs—the online data sets (OLDSs) and write-ahead data sets (WADSs)—shared by the active and the alternate IMS systems
- Databases with access paths from the active and the alternate IMS systems
- For XRF complexes that use USERVAR, a Network Control Program (NCP) running in each IBM 37x5 Communication Controller that controls the SNA terminals at boundary nodes

- Remote terminals that can communicate with either of the IMS systems

The active IMS processes the high-priority IMS workload in the same way that it would without XRF, except that IRLM is not required for XRF. The alternate IMS does not process any IMS transactions; it monitors the active IMS for an indication of trouble. The tasks of the alternate IMS are to track the progress of the active IMS and to update its own control blocks so that work can shift quickly from the active to the alternate IMS.

In an XRF takeover, XRF treats the terminals in differently depending on the class you assign to them. You classify terminals in an XRF system as class-1, class-2 or class-3 depending on the following criteria:

- The kind of terminal
- The teleprocessing access method that controls the terminal
- How the terminal connects to the alternate IMS system
- How your system programmer defines the terminal

Class-1 terminals do not appear to lose connection or have their data processing interrupted in an XRF takeover; however, depending on whether you choose to use MNPS or USERVAR for terminal session management, XRF manages class-1 terminal sessions differently.

With the USERVAR method of XRF terminal session management, the alternate IMS continuously maintains an open backup session for class-1 terminals and switches the terminals to that session in the event of a takeover. With the MNPS method, the alternate IMS does not maintain backup sessions; instead, at the time of a takeover, the alternate IMS opens a new instance of the MNPS ACB. Sessions for class-1 terminals are then rerouted through the new MNPS ACB on the alternate IMS. The differences between each method are transparent to the end user.

For all other terminals, a takeover disrupts sessions. For some terminals, IMS immediately tries to reestablish service at a takeover. These terminals are class-2 terminals. For other terminals, such as those terminals that the operator must manually switch to the alternate IMS system at takeover, no XRF support is available at takeover. These are class-3 terminals. All ETO and LU 6.2 devices are class-3. Users at these terminals depend on an operator to reestablish a session with the alternate IMS system.

**Related reading:** For more information on terminal classes, see “Terminals in an XRF complex” on page 633.

All program temporary fixes (PTFs) that must be applied are identified. Both systems must be at the same IMS release. A planned takeover cannot be used to migrate from one IMS release to another.

Although the alternate IMS has XRF-related tasks to perform, the CPC, z/OS, and VTAM in the alternate IMS are available to process other work as well. Plan the workload carefully, keeping in mind that, at takeover, service is degraded for the non-XRF work. For example, if you run TSO in the alternate IMS, z/OS might swap out the TSO user address spaces at takeover.

The operators of both IMS systems in the complex must enlarge their area of concern and responsibility to include the other IMS system in the XRF complex. Whenever either of the IMS systems changes status, both operators should know about it.

Both operators must be aware of the following behavioral characteristics of the IMS systems in an XRF complex:

- The active IMS system processes the IMS workload.
- The alternate IMS system maintains a state of readiness to take over the workload of the active IMS system.
- Both the active CPC and the alternate CPC can have non-IMS work running on them.

### **XRF variations: MNPS and USERVAR**

An XRF complex can manage terminal sessions during a takeover in one of two ways: using VTAM multinode persistent sessions (MNPS) and MNPS ACBs or using a VTAM USERVAR and USERVAR tables. When designing the XRF complex, you must choose between using MNPS and using USERVAR to manage session switching.

For a terminal user, there is no apparent difference between an XRF complex that uses MNPS and an XRF complex that uses USERVAR; however, each type of XRF complex manages terminal session switching differently during a takeover and an XRF complex that uses USERVAR requires hardware and software that an XRF complex that uses MNPS does not require. Other than the way they manage terminal sessions, the two types of XRF complexes are functionally the same.

The hardware and software requirements of an XRF complex that uses USERVAR include a 37x5 Communication Controller running NCP. An XRF complex that uses MNPS does not require additional hardware or software.

An XRF complex that uses USERVAR requires you to define logon messages, a USERVAR, and APPLID names for session switching to work. An XRF complex that uses MNPS requires the definition of an MNPS ACB on each IMS system in the XRF complex. Terminal users logon using the MNPS name.

An XRF complex that uses MNPS does not maintain backup sessions. In XRF complexes that use USERVAR, the alternate IMS system maintains backup sessions for all sessions that are opened on the active IMS system for terminals defined as class-1.

In an XRF complex that uses MNPS, the alternate IMS system creates no backup sessions. Instead, when a takeover occurs, VTAM maintains persistent sessions for all terminals defined as class-1 that have open sessions on the active IMS. After the alternate IMS system opens a new MNPS ACB, VTAM routes the class-1 sessions through the new MNPS ACB.

For more information about specifying each type of XRF complex, see “Preparing the system for XRF” on page 645.

### **Backup sessions in an XRF complex that uses USERVAR**

In an XRF complex that uses USERVAR, the alternate IMS system opens backup sessions for all class-1 terminal sessions that are open on the active IMS system, thus reducing the time necessary to switch the terminals during a takeover.

Whenever a class-1 terminal user logs on to XRF IMS, the active IMS system records this new session in the log. The alternate IMS system reads this record and initializes a backup session for the terminal. Data does not flow on the backup session as long as it remains a backup session. When the user logs off from the active IMS, the alternate IMS system terminates the backup session.

At a takeover, the alternate IMS requests, through VTAM, that NCP switch each of the class-1 terminals that are communicating with the active IMS from primary sessions to the preopened backup sessions. NCP responds to these requests. Transparency of takeover depends on what is happening on the terminal at the moment of takeover. The user might be unaware of the takeover or might be asked to reenter the last transaction.

## **DBCTL capabilities**

A DBCTL environment alone does not have full XRF capability. It does, however, allow you to have a standby, alternate DBCTL environment.

This alternate IMS does not track the active IMS. You can only bring it to a point where it is fully initialized and ready to take over. The alternate IMS is then called a pre-initialized DBCTL environment. If the active IMS fails, you can send a restart command to the alternate IMS.

A DB/DC environment can provide DBCTL services to a CCTL because the DBCTL services are built into IMS DB. If that DB/DC environment runs XRF, the DBCTL service provided is fully XRF capable. Either the operator or the CCTL can issue a switch command to start the takeover.

---

## **XRF takeover**

The best way to describe a takeover is through an example. Consider a typical installation with a need for high availability: a manufacturing control center that processes production data coming in from many remote terminals. Processing includes updating the IMS databases and recording activity on the IMS system log. Accompanying tasks are running payroll, producing a company report, and performing routine maintenance, which are often lower in priority. This installation runs z/OS with VTAM as the teleprocessing access method.

The failure in this example is an IMS control region abend.

### **IMS failure without XRF**

In an installation with one IMS system and without XRF, the IMS control region abend disrupts all service at the terminals. If the system is unavailable long enough, production lines shut down. The operator, the system programmer, or both might have to:

1. Wait for a diagnostic dump
2. Determine the cause of the abend
3. Restart IMS
4. Restart the associated dependent regions
5. Reestablish sessions with terminals

These activities can take from one hour to several hours depending on the type of outage and the size of the installation. Another half-hour might be required to recover IMS to where it was before the outage. Although the original problem



might have been a minor one, the recovery is costly. End users can be without service for almost two hours while the operational staff performs the recovery steps. Regaining processing ability and ensuring the integrity of the databases and message queues during the down period is necessary for this installation.

## IMS failure with XRF

An installation with XRF installed is physically different from the installation just described. The following figure shows the two systems at the XRF complex. Each system includes the CPC and three licensed programs: z/OS (including DFP), IMS, and VTAM.




Figure 59. The XRF complex before a takeover

IMS 1 processes the high-priority work—production data that comes in from the remote terminals. It updates the databases and records its activity on the IMS system log.

IMS 2 tracks the active subsystem by monitoring the records on the IMS system log. To maintain an environment identical to that in the active IMS, IMS 2 updates many control blocks and message queues in the alternate IMS to reflect those in the active IMS. If your XRF complex uses USERVAR, IMS 2 also opens backup sessions for class-1 terminal users that log onto the active IMS. CPC capacity and storage not used by this activity support the company report program.

When IMS 1 abends, the takeover begins. Depending on the amount of real storage that the XRF processes demand, operating system 2 might swap out the company report program. IMS 2 shifts the production workload to the alternate IMS and, if your XRF complex uses MNPS, opens a new instance of the MNPS ACB. At this point, IMS 2 begins to serve class-1 and class-2 terminals and you can begin determining the problem on the failed IMS 1.

While IMS 2 recovers data and switches sessions on class-1 terminals, IMS 2 and operating system 1 prevent IMS 1 from writing to the IMS system log and the databases. IMS 2 isolates the log and proceeds with the takeover. At the same time, operating system 1 performs I/O prevention, ensuring that all new I/O requests to the databases from IMS 1 return without being executed. When z/OS has completed or canceled all existing I/O requests to the database data sets, it notifies the operator. The following figure shows the two systems during the takeover.




Figure 60. The XRF complex during a takeover

The takeover is complete when all the users at class-1 terminals can communicate with IMS 2 and can enter transactions and receive replies from their IMS applications. When IMS 2 learns that the failing IMS 1 cannot write to the databases, it stops protecting them. The following figure shows the XRF complex after the takeover.






Figure 61. The XRF complex after a takeover

XRF does not perform problem determination for the installation. It does, however, reduce the pressure that exists in a non-XRF installation while the terminal users are without service during dumping and restarting. The installation personnel still must determine the cause of the problem in the failed active subsystem. When dumping activity completes on the failed active subsystem, an operator can return operating system 1, IMS 1, VTAM 1, and CPC 1 to service as the alternate IMS for operating system 2, IMS 2, VTAM 2, and CPC 2. If the dumping activity is lengthy, the operator might bring up a third system as an alternate. Until an alternate IMS exists for the system that is processing the requests of the IMS users, no XRF complex exists.

During the takeover, IMS and z/OS send messages to their system operators to notify them of the progress of the takeover. Depending on the cause of the takeover and the resources at the installation, the operators have some tasks to perform.

#### **Related reading:**

- For information on how z/OS, IMS, VTAM, and NCP contribute to the XRF takeover process, see “Component roles in the XRF process” on page 597.
- For information on the messages that XRF sends to the operators at takeover and the appropriate operator responses, see “Post-takeover phase of the XRF process” on page 621.

### **Takeover conditions**

A takeover condition is an event that causes IMS in the alternate IMS to request a takeover. In an installation that runs XRF, an IMS control region abend is always a takeover condition. Parameters that the system programmers specify when they tailor the IMS determine whether any of the following failures are also takeover conditions:

- A z/OS failure, loop, or wait state

- A CPC failure
- A VTAM failure
- An IRLM failure

**Recommendation:** To prevent an unwanted takeover while IMS processes an SVC dump, ensure that the following conditions exist:

- The XRF takeover interval is sufficient to allow the SVC dump to complete.
- The system data set contains enough space to capture data from the SVC dump.

Not all failures at an IMS installation qualify as takeover situations. XRF does not address the outages caused by failures of service elements you do not duplicate. For instance, XRF does not respond to:

- A channel or link failure that causes a break in communication between the CPC and the communication controllers or DASDs
- Failures in the telecommunication network, such as communication controllers, NCP, lines, and terminals
- Intersystem failures, such as those caused by JES3 or CTCs
- Loss of or damage to the IMS databases
- A power failure that affects both CPCs in the complex
- Failures of user catalogs that point to data sets, such as databases

Each installation must make some planning decisions about the takeover. You decide some of the conditions that initiate a takeover, the amount of operator involvement, which terminal sessions recover automatically at takeover, and which methods the alternate IMS uses to learn of problems in the active IMS.

## Planned takeover

Your installation can take advantage of a takeover to schedule certain kinds of changes to the two systems. While a system is an alternate, your operator can bring it down, leaving the active IMS without backup support for the short time it takes your system programmer to perform the updates. The operator then restarts the alternate IMS with the code or hardware maintenance in place or IPLs z/OS with configuration changes in place. To make the changes on the active IMS, the operator initiates a takeover and repeats the process on the former active IMS.

---

## XRF requirements

Building an XRF complex requires additional processing, storage, and communication resources.

How many additional resources depends on:

- What resources your installation already has
- How heavily utilized your CPCs are

The following list describes the software, hardware, and operational requirements for XRF. For information about the required levels for elements, see *IMS Version 11 Release Planning*.

- XRF licensed program requirements

To use XRF, IMS must have the same system definition in the active and the alternate IMS systems. Before XRF can be operational, both the active and alternate IMS systems must have the required levels of z/OS, DFSMS, SSP, and

VTAM installed (although these licensed programs do not need to be installed at the same time). If your XRF complex uses USERVAR, you must have NCP installed as well.

- XRF hardware requirements

XRF works with the following IBM hardware products:

- CPCs

XRF runs on CPCs supported by z/OS. The CPCs can be different models.

- 37x5 Communication Controllers

If your XRF complex uses USERVAR, XRF requires that you connect the class-1 terminals, class-2 terminals, and intermediate routing nodes (IRNs) to 37x5 Communication Controllers to automatically switch sessions at takeover.

- Terminals

All the terminals in your installation that use IMS can operate in an XRF complex. The level of support varies, as described in “Overview of XRF” on page 585.

Terminals using the X.25 communications protocol can be used with IMS. X.25 terminals can have class-1, class-2, or class-3 XRF support, based on their definition in IMS and their X.25 characteristics. X.25 NCP Packet Switching Interface (NPSI) participates in XRF terminal switching for eligible SNA terminals that are in session with IMS. Eligible SNA terminals attach to X.25 networks with the following:

- Integrated X.25 adapters
- Network Interface Adapter (5973-L02)
- External SDLC/HDLC PAD

Only eligible SNA terminals can receive class-1 XRF support.

For more information on attaching X.25 terminals, see:

- *X.25 NCP Packet Switching Interface General Information Version 3*
- *NCP X.25 Planning and Installation*

XRF terminal switching is limited to SNA terminals. XRF does not support non-SNA devices that communicate with SNA hosts through other NPSI functions, including:

- The protocol conversion for non-SNA equipment (PCNE)
- Packet assembly/disassembly (PAD)
- General access to X.25 transport extension (GATE)
- Dedicated access to X.25 transport extension (DATE)

- XRF operational and management requirements

Although additional personnel are not needed to support XRF, additional tasks are required for system programmers and operators. The z/OS, IMS, and network operators must understand the XRF process, be alert to the possibility of a takeover, and be ready to communicate with each other at takeover. You might need to organize your system consoles so that the operators responsible for managing the active IMS can easily communicate with the operators of the alternate IMS.

The system programmer must initialize z/OS, IMS, VTAM, and NCP by using XRF-related parameters and macros. If your user applications conform to IMS standards, they run in the XRF complex and receive the availability improvements offered by XRF. System programmers can use standard diagnostic aids in an XRF complex.

Before installing XRF at your installation, you must prepare your data sets for minimum disruption during takeover. Ensure that:

- All required data (duplicate z/OS data sets and nonvolatile application and subsystem data) is duplicated on active and alternate systems.
- Access to application databases and logs is through DASD that is shared by the active and the alternate systems.

If you decide to keep the nonresident ACBs in 64-bit storage, the XRF standby system needs to be set up in the same manner by ensuring that the size of the 64-bit storage pool in the standby system is the same as the size of the 64-bit storage pool in the active system.

XRF does place additional demands on your installation's resources. However, if you understand how XRF works and plan carefully, you can reduce this overhead. The effort can be worthwhile if your installation is seriously affected by lengthy planned or unplanned interruptions of IMS service.

- **IMSplex requirements for XRF**

In an IMSplex with the Common Service Layer (CSL), you must define the alternate IMS system with access to CSL. The XRF alternate system must be defined with a unique IMSID in the IMSplex. A Structured Call Interface (SCI) must be defined on the operating system of the alternate IMS system. To route commands to the alternate and active IMS systems with Operations Manager (OM), use the IMSID on the ROUTE() parameter. After a takeover, the new active IMS system takes over ownership of resources defined to the Resource Manager (RM).

If you use the online change function in an IMSplex that does not have RM, the active IMS system must exclusively own the OLCSTAT data set. The OLCSTAT data set can include the IMSID of only the active IMS system.

When the alternate IMS system becomes the active IMS system after a takeover, you must change the IMSID in the OLCSTAT data set to match the IMSID of the alternate IMS system before you can issue an INITIATE OLC command. You can change the IMSID in the OLCSTAT data set by using the Online Change utility (DFSUOLC0).

**Related concepts:**

“System support for application programs” on page 31

---

## Component roles in the XRF process

IMS, z/OS, and several z/OS elements work together in the XRF process.

Being familiar with the functions that each perform will make it easier for you to plan, install, and operate XRF, as well as diagnose any problems that occur.

## Contribution of IMS to the XRF process

IMS is the major contributor to the XRF process, just as the IMS user is the greatest receiver of XRF service. While most of the IMS processing in the active IMS is serving the requests of the IMS terminal users, all of the IMS processing in the alternate system is related to XRF. The alternate IMS system refuses to accept logon requests from users.

The following list summarizes the activities of the two IMS systems before and during a takeover.

- The active and alternate IMS establish and maintain constant communication through the IMS system log and through simple signals that the alternate IMS receives from active IMS. Surveillance mechanisms provide the signals.
- The active IMS:

- Processes the requests of the IMS users.
- Sends surveillance signals across the ISC link and the restart data set (RDS).
- Opens and closes sessions for terminals.
- Opens and closes databases.
- Reports its activities in the log.
- The alternate IMS:
  - Responds to normal startup procedures, but does not process transactions until after a takeover.
  - During normal operations:
    - Uses information from the log to update its control blocks.
    - Allocates and opens the same databases that the active IMS opens.
    - If the XRF complex uses USERVAR, opens and closes backup sessions for class-1 terminals that log on to and log off from XRF IMS.
    - Through surveillance mechanisms, monitors the active IMS and initiates takeover if surveillance indicates that useful work in the active IMS has ceased.
  - During the takeover:
    - Participates with the other licensed programs in taking over the workload of the active IMS. It obtains the workload by taking control of the IMS system log.
    - Prevents the failing active IMS from writing to the log.
    - Recovers the in-flight transactions (those transactions that the active IMS only partially processed).
    - If the XRF complex uses MNPS, opens a duplicate instance of the MNPS ACB.
    - Starts executing new transactions.
    - Starts session recovery on class-1 and class-2 terminals. “Session recovery” in this information refers to IMS attempts to switch sessions on class-1 terminals or to reestablish service on class-2 terminals.

## IMS as a recoverable service element

The two IMS systems in the XRF complex have a special status in the complex. Although failures in IMS, z/OS, VTAM, and the CPC can cause takeovers and are therefore recoverable, XRF considers only IMS the recoverable service element (RSE). z/OS, VTAM, and the CPC are dependent service elements (DSEs); they depend on IMS to recognize their failures and then to request a takeover. The RSE name appears in messages that the operators receive at takeover. For example, two IMS systems, named IMS1 and IMS2, form the RSE named IMSPROD. After a takeover at this complex, a z/OS message tells the operators that SUBSYSTEM IMS2 IS NOW THE ACTIVE ELEMENT OF RSE IMSPROD (message AVM007I).

**Related reading:** For more information about the RSE name, see RSENAME in “Defining IMS.PROCLIB members for XRF” on page 649.

## Surveillance mechanisms

During normal operations, the alternate IMS tracks the active IMS by reading the IMS system log. The active IMS writes its activities on the log, and the alternate

system checks these records, updating its own control blocks. In this way, control blocks in the alternate system reflect those in the active IMS, and the alternate IMS remains ready to take over.

The active IMS records on the system log the following failures:

- IMS control region abend
- VTAM failures that lead to TPEND exits in IMS
- IRLM failures that lead to STATUS exits in IMS

Certain error conditions can prevent IMS from recording these failures on the system log. An example of such a condition is the inability of IMS to purge the OLDS buffer.

The log does not warn of other failures. For this reason, surveillance mechanisms provide periodic signals from the active IMS. The alternate IMS monitors these signals. When the signals fail to appear, the alternate IMS considers a takeover. Absence of these signals can indicate:

- z/OS failures
- z/OS loops or wait states
- CPC failures

The alternate IMS can receive these signals in three ways as illustrated in the following figure:

1. The active IMS sends messages over the ISC link between the active and the alternate IMS systems.
2. The active IMS places a time stamp in the RDS.
3. The active IMS continues to add new records to the IMS system log.

Using the parameters in member DFSHSBxx of IMS.PROCLIB, you choose which methods of surveillance operate in your installation. The following figure illustrates the three surveillance mechanisms. The ISC link can be a VTAM CTC or a shared line through the 37x5 Communication Controller.




Figure 62. Surveillance options

The signals across the ISC link and on the RDS are periodic. To use the IMS system log as surveillance, the alternate IMS periodically checks for new records on the log.

A failure to receive the signals on the ISC link and the RDS data set can cause a takeover; an absence of new records on the log cannot. However, if the alternate IMS knows of log activity, it overrides a takeover indication from the RDS or the ISC link.

In parameters you define in the DFSHSBxx member of IMS.PROCLIB, you specify:

- Which surveillance mechanisms you want your installation to use
- How often the active IMS is to send signals on the ISC link or the RDS
- How often the alternate IMS is to check the log for a new record
- Which absences of signals within specified times are takeover conditions

“Establishing surveillance for XRF” on page 606 describes how to code these parameters. Your operator can stop and start surveillance dynamically by using IMS /STOP and /START commands. The operator can use the /CHANGE command to change the surveillance timing intervals.

## Contributions of z/OS and the z/OS elements

z/OS components and other products actively participate in XRF processing. Some of these components and products are the Availability Manager (AVM), Data Facility Storage Management Subsystem (DFSMS), Virtual Telecommunications Access Method (VTAM), Network Control Program (NCP), System Support Program (SSP), and Tivoli NetView for z/OS.

### Availability Manager

Through its availability manager (AVM) element, z/OS provides the environment for the active and the alternate IMS and provides services during a takeover. Specifically, AVM does the following:

- Provides the I/O prevention that keeps the failing IMS system from accessing the databases
- Sends AVM-prefixed messages describing the status of the XRF complex to operators during the takeover

These messages have the prefix, AVM. They describe the status of the XRF complex and help the operators respond correctly to the takeover.

The system resource manager (SRM) component of z/OS hastens the acceptance by the alternate IMS of the new workload when a takeover begins. At this time, the alternate IMS temporarily requires more real storage to recover databases and sessions. The SRM component of z/OS speeds up its analysis of the need for storage by the alternate IMS. This frequent checking allows SRM to respond quickly to this need.

*I/O prevention* is the action AVM takes to stop the failing IMS from writing to the databases. It begins when the active z/OS learns that IMS is requesting a takeover. At this time, the active IMS might have scheduled updates to the databases. To maintain database integrity, AVM in the active z/OS ensures that current I/O operations are complete and that z/OS does not honor any additional I/O requests to the databases from the active IMS. When AVM is sure the database data sets are safe, it issues a message that announces I/O PREVENTION IS COMPLETE (message AVM006E).



Of course, z/OS cannot prevent I/O if it is no longer operating. If z/OS cannot prevent I/O, an operator of the failing active z/OS must manually make sure that the active IMS system does not write to the databases.

**Related reading:** For a description of the operational procedures, see *IMS Version 11 Operations and Automation*.

When an operator at the failing active IMS system is certain that IMS is no longer changing the databases, that operator informs an operator at the alternate IMS system. Then the operator at the alternate IMS responds "GO" to the message that says REPLY "GO" WHEN I/O PREVENTION COMPLETES (message AVM006E). This action (or the /UNLOCK SYSTEM command) completes the efforts to ensure the integrity of the databases by the alternate IMS.

Your system programmers and operators must understand the importance of preventing IMS in the failing active IMS from accessing the databases. The integrity of the databases is lost if both IMS systems can write to them at the same time.

**Related reading** For information about establishing takeover procedures, see *IMS Version 11 Operations and Automation*.

## **Data Facility Storage Management Subsystem (DFSMS)**

XRF needs DFSMS for I/O prevention, VSAM, and media manager. DFSMS, however, does not produce any messages or change any existing procedures.

## **Virtual Telecommunications Access Method (VTAM)**

The contributions of VTAM differ slightly depending on whether your XRF complex uses MNPS or USERVAR; however, for terminal users, there is no apparent difference.

The VTAM contributions to the XRF process include:

- VTAM allows the user to log on to the IMS systems in the XRF complex using a single logon message. The user has no need to know which one of the IMS systems is currently active.

With XRF, the two IMS systems in the complex appear to terminal users as a single IMS system. However, to VTAM they are unique applications. VTAM allows a terminal user to log on to XRF IMS with a logon message that you choose or, a terminal user can log on using the command LOGON APPLID that specifies the logon message.

- VTAM enables the alternate IMS system to maintain sessions for class-1 terminals that are logged on the active IMS in the event of a takeover. In an XRF complex that uses MNPS, VTAM reroutes sessions for the class-1 terminals through a new instance of the MNPS ACB on the alternate IMS system. In an XRF complex that uses USERVAR, VTAM enables NCP to switch class-1 terminals to their backup sessions on the alternate IMS system.

VTAM operations in an XRF complex differ depending on whether USERVAR or MNPS is used.

### **VTAM operations in an XRF complex with USERVAR**

VTAM checks its interpret table for the variable that corresponds to the logon message. VTAM then looks in the USERVAR table for the VTAM application name that corresponds to this variable. This application name



identifies the IMS system for the terminal to communicate with at this time. This IMS system is the session partner for the terminal.

Be sure that you understand the two tables and the entries in the tables. Each table contains two columns. The interpret table has a logon message column and USERVAR variable column. The USERVAR table has a USERVAR variable column and a VTAM application name column. VTAM uses the USERVAR variable to associate the logon message with the application name of the session partner. It resides in both the interpret table and the USERVAR table in the USERVAR variable column. The following figure illustrates the interpret table and the USERVAR table. The two entries in the logon message column of the interpret table allow users to log on to XRF IMS with the logon messages IMSP or IMSA.




Figure 63. The interpret and USERVAR tables

The interpret table is a static table that you build at VTAM initialization. You use the LOGCHAR macro to place entries in this table. Entries in the USERVAR table are established at initialization but can be changed dynamically. IMS issues the MODIFY USERVAR command to initially place an entry in the USERVAR table. An operator or an application program uses the same command to delete a user-managed USERVAR and specify that VTAM manage the USERVAR automatically, or to change an entry in the USERVAR table. For example, at takeover, the MODIFY USERVAR command changes the application name in the USERVAR table to reflect the new session partner for all terminals.

**Related reading:** For more information on the VTAM USERVAR table, see “Defining the VTAM USERVAR table” on page 656.

The top part of the following figure shows the interpret and USERVAR tables for the VTAM that owns the terminals. In this case, it is VTAM in the alternate IMS. When a user logs on with the logon message IMSP, VTAM searches the interpret table for IMSP and finds the corresponding USERVAR IMS. VTAM then searches the USERVAR table for the application name that corresponds to IMS and opens the session with the IMS system IMS1.

The bottom part of the following figure shows that, following a takeover, VTAM2, at the request of IMS2, has changed the application name in the USERVAR table. Now when a user at a terminal specifies IMSP, VTAM connects the user to IMS2.




Figure 64. VTAM processing of the logon in an XRF complex that uses USERVAR

All existing user programs that conform to present IMS standards continue to execute in the XRF complex. If the USERVAR is VTAM-managed, a modification to the new VTAM APPLID is done automatically. If you have programs still using the logon that they used previously and that are not VTAM-managed, you must make a modification, because the logon does not match the VTAM APPLID of the active IMS.

To modify application programs with user-managed USERVARs, do the following:

- Use INQUIRE USERVAR to communicate with the active IMS to determine the real name of the current active IMS.
- For terminal programs sensitive to the PLU/SLU name in BIND data, the USERVAR is appended by IMS in BIND data.
- New VTAM SENSE codes need to be tested.

IMS prevents a user from logging on to the alternate IMS with the message INVALID LOGON REQUEST IN THE BACKUP SYSTEM (message 3862I). Only IMS master and secondary terminals, the system console, and the ISC surveillance link can communicate with the alternate IMS. The term BACKUP in this message refers to the alternate IMS in the XRF complex. In IMS messages, the term BACKUP can also refer to the backup sessions on class-1 terminals.

#### VTAM operations in an XRF complex that uses MNPS

In an XRF complex that uses MNPS, terminal users log on by using the MNPS ACB name. This name is the same for the MNPS ACBs of both the active and alternate IMS; however, only one MNPS ACB exists at a time. Prior to a takeover, only the MNPS ACB of the active IMS exists. After takeover, only that of the alternate exists. Regardless of which MNPS ACB exists, the logon information that terminal users enter remains the same.

VTAM routes all terminal sessions through the MNPS ACB. In the event of a takeover, VTAM maintains as persistent all sessions for class-1 terminals until the MNPS ACB of the active IMS is shut down and the MNPS ACB of the alternate IMS is opened. After MNPS ACB on the alternate IMS is opened, VTAM routes all terminal sessions through this new ACB.

The following figure shows an example of VTAM ownership of a terminal session in an XRF complex that uses MNPS and how VTAM manages the session prior to a takeover. In this case, VTAM 1 owns the active IMS on CPC 1. Logging on to CPC 2, VTAM 2 routes the terminal's session to VTAM 1, which connects it to the MNPS ACB and the active IMS.




Figure 65. VTAM processing of the logon in an XRF complex that uses MNPS, part 1

The following figure illustrates how VTAM reroutes the terminal session after a takeover. Here, a new instance of the MNPS ACB has been opened on CPC 2 by the alternate IMS. VTAM 2 now connects the terminal session directly to the MNPS ACB without passing the session to VTAM 1, even if VTAM 1 continues to run. The terminal user is unaware of this switch.




Figure 66. VTAM processing of the logon in an XRF complex that uses MNPS, part 2

In z/OS version 1.6 and above, XRF complexes that use MNPS take advantage of the VTAM persistent session forced takeover support. Forced takeover support allows the alternate IMS system to take over when the active IMS has failed, but still appears active to VTAM.

The active IMS system allows a forced takeover by executing the VTAM SETLOGON OPTCD=PERSIST macro with FORCETKO specified. VTAM then allows the alternate IMS system to open the MNPS ACB during XRF takeover processing.

## Network Control Program (NCP)

In an XRF complex that uses USERVAR, NCP maintains the control blocks for the backup sessions and performs the switching of sessions when the alternate IMS requests it. NCP does not have a role in an XRF complex that uses MNPS.

## System Support Program (SSP)

In an XRF complex that uses USERVAR, the SSP defines and generates an NCP. The SSP also loads the program into the 37x5 Communication Controller and dumps the contents of the 37x5 Communication Controller back to the host CPC for diagnostic purposes. SSP does not produce any messages or change any existing procedures. The SSP does not have a role in an XRF complex that uses MNPS.

## Tivoli NetView for z/OS with USERVAR

Your installation probably uses the services of the licensed program Tivoli NetView for z/OS. Although not modified for XRF, Tivoli NetView for z/OS can signal a change in the VTAM application name to VTAMs inside and outside the XRF complex. If a VTAM in the network is not at Version 4 with the USERVAR management enhancement, an Tivoli NetView for z/OS CLIST should be used to propagate a USERVAR change to that VTAM. This change in the VTAM application name occurs at initialization—at completion of IMS restart or processing of the /START DC command—or at a takeover. If your installation does not have Tivoli NetView for z/OS, or if a VTAM is below Version 4, operators in each system in the network must issue the MODIFY command to update the VTAM application name.

### Establishing surveillance for XRF

Use the following five parameters to establish surveillance for XRF: SURV, LNK, LOG, RDS, and SWITCH.

The following table describes the planning decisions, the corresponding parameters, and the options on the parameters. The surveillance parameter and options on the parameters refer to:

- LNK as the ISC link between the two systems
- LOG as the IMS system log
- Restart data set (RDS) as the RDS

*Table 46. Surveillance options on parameters*

| Planning Decision                                                                                                            | Parameter | Options       |
|------------------------------------------------------------------------------------------------------------------------------|-----------|---------------|
| What surveillance mechanisms should operate in the XRF complex? (Step 1)                                                     | SURV      | LNK, LOG, RDS |
| If you specified LNK as surveillance, how long is the interval between signals? (Step 2)                                     | LNK       | Interval      |
| If LNK is a takeover condition, how long should the alternate IMS wait for a signal before considering taking over? (Step 4) | LNK       | Timeout       |
| If you specified LOG as surveillance, how long is the interval between signals? (Step 2)                                     | LOG       | Interval      |
| If LOG is a takeover condition, how long should the alternate IMS wait for a signal before considering taking over? (Step 4) | LOG       | Timeout       |
| If you specified RDS as surveillance, how long is the interval between signals? (Step 2)                                     | RDS       | Interval      |
| If RDS is a takeover condition, how long should the alternate IMS wait for a signal before considering taking over? (Step 4) | RDS       | Timeout       |
| Should the absence of specific surveillance signals cause the alternate IMS to request a takeover? (Step 3)                  | SWITCH    | LNK,LOG,RDS   |

To establish surveillance, perform the following four steps:

1. Determine which surveillance mechanisms to use for the complex.

Use the SURV parameter to establish the surveillance mechanisms. You can specify combinations of the three mechanisms: LNK, LOG, and RDS. Because of the importance of surveillance, however, you should use all three mechanisms.

The alternate IMS uses these three surveillance mechanisms in different ways. Failure of the alternate IMS to receive signals on the ISC link or the RDS is possible cause for takeover; failure of the alternate IMS to receive new records on the log is not. Rather, IMS uses the LOG option to confirm or override a decision that the alternate IMS might make based on lack of signals from the RDS or the ISC link. For example, suppose you specify SURV=(LNK,LOG), and the alternate IMS stops receiving signals on the ISC link. The failure of signals over the ISC link indicates a takeover, but as long as the alternate IMS continues to receive log records satisfactorily, it does not request a takeover. In this case, IMS assumes that the ISC link itself is experiencing difficulty and is therefore not a reliable indicator.

In an XRF complex that uses USERVAR, if you do not have an ISC link already in place, establish it through the 37x5 Communications Controller that controls your class-1 terminals.

The surveillance mechanisms and the internal parameters can be changed dynamically by commands issued from the master terminal.

**Related reading:** For more information on surveillance mechanisms, see “XRF parameters in DFHSBxx” on page 650.

2. For each surveillance mechanism you choose, decide the interval value (how often the alternate IMS is to receive the signals).

After you have established which surveillance mechanisms are to operate in the complex, you should decide how often the alternate IMS should receive signals through the surveillance.

Code the LNK and RDS parameters to establish the intervals between signals for the LNK and the RDS for each of the surveillance mechanisms you choose.

Two factors to consider in setting the timing for LNK and RDS are:

- The speed of communication over the ISC link.
- The performance overhead on the alternate z/OS operating system. Set the timing interval higher if you need to reduce the work of the alternate z/OS operating system.

Code the LOG parameter to establish how often the alternate IMS should check the log for new records.

For more information about interval values, see the topic “XRF parameters in DFHSBxx” on page 650.

**Related reading:** IMS has specific rules and recommendations for coding the interval values for the two systems. See *IMS Version 11 System Definition* for these rules.

3. Decide whether the absence of any or all of these signals should be a takeover condition.

Use the SWITCH parameter to establish absence of surveillance signals as criteria for takeover. Specify on the SWITCH parameter the surveillance mechanisms that are critical in alerting the alternate IMS of certain events. For example, an installation managing its local locks on the active IMS with IRLM, or one using data sharing, would specify the IRLM failure as a criterion for takeover. Installations with a large VTAM network would probably specify a takeover to occur when a VTAM failure causes an IMS TPEND exit. An installation with a large non-VTAM network and a small VTAM network

would probably not specify a VTAM failure as a criterion for takeover. For more information about the SWITCH parameter, see the topic “XRF parameters in DFSHSBxx” on page 650.

4. If the answer to 3 on page 607 is yes, decide the timeout value (the length of time the alternate IMS waits for a surveillance signal before considering a takeover).

After you have specified which surveillance signal absences are takeover conditions, you can decide how long the alternate IMS should wait for a signal before requesting a takeover. Specify timeout values on the LNK, LOG, and RDS parameters. The default intervals are 9 seconds for the LNK parameter and 3 seconds for LOG and RDS.

If timeout values are too low, some z/OS recovery routines that are transparent in a non-XRF environment might cause takeovers. An example of such a routine is an alternate CPU recovery (ACR) routine. These unwanted takeovers are more likely to occur on a larger CPC or with a large IMS workload.

---

## Phases of the XRF process

The two systems that form your XRF complex pass through six identifiable phases.

The phases of XRF are:

- Initialization
- Synchronization
- Tracking
- Takeover
- Post-takeover
- Termination

You can determine the current phase at any time by checking the system status line on the console for the alternate IMS.

Each subsystem participates in each phase, but the active IMS has little XRF activity in the synchronization, tracking, and takeover phases. The subtopics in of this topic describe the activities for the active and alternate subsystems for each of the six phases.

### Initialization phase of the XRF process

The initialization phase begins when an operator brings up AVM as a z/OS started task, either through the START AVM command or through the /NRESTART or /ERESTART command.

The START command loads modules and control blocks and, optionally, opens a session on the ISC link that connects the active and the alternate IMS systems.

The following figure shows the XRF initialization phase.




Figure 67. XRF initialization phase

An operator at the active IMS uses z/OS START commands to bring up the active IMS, as in a normal or emergency restart. When it is up, the operator at the alternate IMS can bring up the alternate IMS at any time. The commands are similar for both systems except that the operator at the alternate IMS adds the option AUTO=N to the START IMS command and issues the IMS command, /ERESTART BACKUP.

**Attention:** An operator error can result in losing system or database integrity.

The alternate IMS is in a state similar to emergency restart until a takeover occurs. To be ready for a takeover, the alternate IMS shares the databases with the active IMS. The alternate IMS opens the databases, but does not access them until after a takeover.

In addition to the normal IMS initialization functions, IMS does the following to set up the XRF complex:

1. Processes the new DFSHSBxx IMS PROCLIB member. For more information, see “XRF parameters in DFSHSBxx” on page 650.
2. Loads the new XRF/IMS modules, and obtains the new pools and work areas. Most of the new modules reside in the control region private below the 16 MB line.
3. Identifies both subsystems to the z/OS availability manager using the RSENAME keyword, described in “Defining IMS.PROCLIB members for XRF” on page 649.
4. Establishes the DBRC environment on both the active and the alternate subsystems using the RSE name. If a takeover occurs, authorizations of the active IMS are automatically passed to the alternate IMS.
5. Informs VTAM of the APPLID of the active IMS system.

**Related reading** For more information on how the operators bring up the active and alternate IMS systems, see *IMS Version 11 Operations and Automation*.



## Synchronization phase of the XRF process

During the synchronization phase, the alternate IMS becomes an image of the active IMS, as shown in the following figure. The phase begins when the environments are fully initialized.




Figure 68. XRF synchronization phase

The process of building the control blocks in the alternate IMS is very simple. The alternate IMS determines if the last checkpoint was a SNAPQ. If it was not, and you have an ISC link, the alternate IMS requests that the active IMS take a SNAPQ checkpoint. The SNAPQ log records are used to check that the IMS system definitions on both systems match and to build an initial “snapshot” of the active IMS on the alternate IMS.

If you do not have an ISC link, the alternate IMS issues a message to ask the master terminal operator at the active IMS to request the SNAPQ checkpoint.

The alternate IMS opens the IMS system log, locates the SNAPQ checkpoint, and reads from the log data sets.

If the ISC link exists, a message is sent to the alternate IMS while the checkpoint is being written to the log. If the ISC link does not exist, the alternate IMS monitors the restart data set (RDS) for the completion of SNAPQ.

IMS uses the contents of the dump for several things:

- The alternate IMS determines which MODSTAT and RDS are being used by the active IMS.
- The alternate IMS tries to preallocate and pre-open databases and areas that are accessed by the active IMS.

- MSDBs are loaded by the alternate IMS from the SNAPQ checkpoint records; separate MSDB checkpoint data sets must be available for the alternate IMS.
- If your XRF complex uses USERVAR, backup sessions for terminals that are opened on the active IMS are initiated on the alternate IMS for those class-1 terminals that are defined with the BACKUP option.
- Security control blocks are loaded on the alternate IMS from the RACF data set if you are using RACF for security.
- The DBRC subsystem record is updated to indicate that an alternate subsystem exists. This suppresses the DBRC 'SIGNOFF ABNORMAL' message in the case of failure of the active IMS.
- Status of dependent regions on the active IMS is tracked on the alternate IMS.

Only after a takeover can the alternate IMS write to the log.

Operators can start dependent regions in the alternate IMS to have them ready for a takeover.

When the alternate IMS has processed the SNAPQ checkpoint, the synchronization phase is complete. When the control blocks of the alternate IMS are synchronized with the ones in the active IMS, the alternate IMS is able to take over the processing of the active IMS. From this point on, you can request a planned takeover.

## Tracking phase of the XRF process

Most of the time, your XRF complex is in the tracking phase. The active IMS processes the requests of IMS users, writes to the log, and sends signals across the surveillance mechanisms that were set up in the initialization phase. The alternate IMS tracks the active IMS to reduce the work that the alternate IMS must do at takeover.

The tracking phase looks similar to the XRF synchronization phase that was shown in "Synchronization phase of the XRF process" on page 610.

Tracking activities of the alternate IMS include:

- Reading the log
- Using log information to update control blocks
- If your XRF complex uses USERVAR, opening and closing backup sessions for class-1 terminals
- Checking the surveillance mechanisms and the log records for signs of failure in the active IMS

Tracking activities use few of the non-storage resources of the alternate IMS. The use of real storage depends on how you page fix IMS storage and whether you define a member of DFSFIXxx with page-fixing options that are in effect only when the subsystem is active.

## Updating control blocks in the alternate IMS

The alternate IMS reads the OLDS (Use dual OLDSS, because tracking can continue using the other OLDS if one OLDS is lost.) to track:

- Communication network status

In addition to the usual terminal status tracking that is done during the emergency restart process, the alternate IMS monitors the terminal connection

status. This is done initially from records in the SNAPQ checkpoint and subsequently from terminal connection status change log records written by the active IMS. For class-1 terminals in an XRF complex that uses USERVAR, backup sessions are maintained by the alternate IMS. When class-1 terminal sessions are terminated on the active IMS, their backup sessions are also terminated on the alternate IMS. The user can never log onto the alternate IMS. Terminal types are described in “Terminals in an XRF complex” on page 633.

- Application program and transaction status

As application programs are scheduled on the active IMS, the alternate IMS builds the control block structure to track the dependent region status.

The operator can prestart dependent regions on the alternate IMS to allow the alternate IMS to complete takeover more quickly.

- Database status

As databases and areas are allocated and opened on the active IMS, the alternate IMS pre-allocates and pre-opens these same databases and data set areas. The IMS macro definitions for all databases and areas must be specified as SHARED. The alternate IMS also tracks all DBRC database authorizations.

**Recommendation:** Register all databases and areas in DBRC. The alternate IMS also tracks the status of all locks for uncommitted database changes.

- Message queue status

Message queues are built on the alternate IMS from the SNAPQ checkpoint records. They are updated on the alternate IMS by OLDS records. A local message queue is used by the alternate IMS to communicate with the operator. During IMS system definition, two master and two secondary master terminals are specified. The master terminals at the active IMS control the active IMS; the master terminals at the alternate IMS control the alternate IMS.

- MFS pool status

MFS blocks are preloaded and released on the alternate IMS as they are loaded and released on the active IMS.

- Dependent region status

The alternate IMS initializes and maintains control blocks to reflect the status of the dependent regions. It monitors all activity in the region so that backout of uncommitted processing can occur at takeover.

With this information, the alternate IMS updates its control blocks to maintain an environment that duplicates the active IMS. It pre-allocates and pre-opens the IMS databases and data areas as the log indicates that the active IMS allocates and opens them. It also loads the MFS control blocks, program specification blocks (PSBs) for backout use, and data management blocks (DMBs).

## Surveillance of the active IMS

As described in “Surveillance mechanisms” on page 598, the alternate IMS depends on surveillance to indicate some of the takeover conditions in the active IMS. If surveillance includes the restart data set (RDS), the active IMS periodically places time stamps on the RDS. If surveillance includes the ISC link, the active IMS sends periodic signals over the ISC link. If surveillance includes the log, the alternate IMS periodically checks to ensure that log records continue to arrive.

## Takeover phase of the XRF process

The alternate IMS requests a takeover for one of three reasons.

- A problem occurs that is so serious that processing in the active IMS cannot continue. These problems include:
  - An IMS abend
  - A z/OS failure, loop, or wait state
  - A CPC failure

These problems are based on the assumption that you establish at least basic surveillance including the restart data set (RDS) and the ISC link. If surveillance detects a failure in the active IMS, or if the operator issues a /SWITCH SYSTEM FORCE command, the alternate IMS reserves the logs to prevent the active IMS from further processing and proceeds with the takeover. The operator must terminate the old active IMS or reset the CPC to prevent further updates to the databases. If the active IMS abends, AVM in the active IMS invokes I/O prevention to prevent further updates to the databases.

- The operator initiates a planned takeover.

To initiate a planned takeover, an IMS operator at the active or alternate IMS issues the /SWITCH SYSTEM command to cause a takeover. IMS waits until all in-flight messages are completed for every dependent region before it abends. AVM in the old active IMS performs I/O prevention, and no back outs have to be performed on the alternate IMS. An example of the use of the planned takeover is when your system programmer wants to do maintenance in the active IMS. “Practical uses of the planned takeover” on page 620 suggests other uses of a planned takeover.

- Other criteria that you have set in the DFSHSBxx PROCLIB member has triggered a takeover.

As mentioned in “Surveillance mechanisms” on page 598, you can specify that an absence of a signal from surveillance causes a takeover. For example, you can tell the alternate IMS that the absence of a signal over the ISC link for ten seconds is a takeover condition.

Two other failures that you can set up as criteria for takeover are:

- Abend of the IRLM

You have IRLMs at your complex for two reasons:

- The IRLM in the active IMS might control access to the databases that XRF IMS shares at a block level with IMS systems outside the complex. In this case, your XRF complex has two IRLMs, one at the active subsystem and one at the alternate subsystem.
- IRLMs might be managing your local locks.

In either case, you might want a takeover to occur if the IRLM in the active IMS terminates, at which point, an IMS STATUS exit routine is invoked.

- Failure of VTAM

You can direct the alternate IMS to request a takeover if VTAM invokes the IMS TPEND exit routine.

Failures of VTAM and IRLM cause degraded performance at your installation, but you might still decide to avoid a takeover if these failures occur.

You specify the absence of surveillance signals or VTAM or IRLM failures as takeover conditions in the parameters you add to member DFSHSBxx of IMS.PROCLIB. “Establishing surveillance for XRF” on page 606 describes these parameters.

The alternate IMS checks the signals it receives on the RDS and the ISC link, reads the records it receives on the log, and monitors the commands of the operators. Eventually it requests a takeover. At this point, the takeover begins. Depending on

how IMS is tailored, the takeover proceeds automatically or waits for an operator to perform certain tasks before allowing the takeover to proceed. The takeover phase is illustrated in the following figure.




Figure 69. XRF takeover phase

During takeover, the alternate IMS uses the data it collected during the synchronization and tracking phases to transform itself into an active IMS. It also ensures that work on the former active subsystem halts. XRF's primary objective during this phase is to ensure the integrity of the databases. At the end of this phase, the alternate IMS is the new active IMS. After the takeover begins, it cannot be revoked—it proceeds to completion.

## The takeover begins

Because the condition of the failing active IMS at takeover is uncertain, the alternate IMS does most of the work at takeover. The alternate z/OS operating system must decide what to do with the non-XRF workload that was processing on the alternate IMS. See “Non-XRF workload on the alternate IMS” on page 632.

The events described in this subtopic occur during a takeover if IMS abends or if an operator at the active IMS issues the `/SWITCH SYSTEM` command. See “Takeover processes without I/O prevention” on page 618 for a description of the takeover process caused by other events.

During a takeover, the active IMS attempts to do the following:

- Recognizes the failure and writes a termination message to the log. It stops processing incoming messages.
- Issues a `RESERVE` against the OLDS and WADS of the active IMS if the active and the alternate subsystems reside on different CPCs.

**Recommendation:** No other data sets should reside on these devices, because they are inaccessible after the `RESERVE` is issued. Access to these data sets when a takeover has begun might also hold up the takeover.

If the active and the alternate IMS systems reside on the same CPC, the activities differ: the alternate IMS must ABTERM the active IMS and wait until the IMS resource clean-up routine on the active IMS notifies the alternate IMS that I/O to the log has terminated on the active IMS.

- Closes the OLDS using the WADS when IMS can no longer access the logs. Because a takeover is not viewed as a failure by DBRC, the alternate IMS switches to the next set of previously allocated OLDS.
- Invokes the z/OS system resource manager (SRM), which tries to obtain resources for the alternate IMS. Because the alternate IMS temporarily requires more real storage to recover databases and sessions, SRM hastens the acceptance of the new workload by the alternate IMS.
- Removes incomplete messages during takeover. These incomplete messages occur at takeover when the message queues have already been built on the alternate IMS.
- Defers the deletion of nonrecoverable messages.
- Re-enqueues recoverable messages that were in process at the time of failure, unless they must be backed out. If backout is required, these messages are automatically re-enqueued at the completion of each backout.
- Merges messages on the local message queue so that the IMS master terminal operator does not lose critical messages due to the takeover.

The active IMS performs the listed tasks except when it cannot purge the OLDS buffer.

If AVM is operating in the active IMS, it performs I/O prevention. Normally, I/O prevention is completed within a few seconds of the I/O prevention request. “Contributions of z/OS and the z/OS elements” on page 600 describes the action AVM takes to stop the failing IMS from writing to the databases. If AVM is not operating, the operator at the active IMS must manually perform I/O prevention by resetting the CPC or terminating IMS. At the completion of I/O prevention, an operator at the active IMS must inform an operator at the alternate IMS of the completion, and the operator at the alternate IMS must reply “GO” to message AVM005A.

Even before this notification, the alternate IMS proceeds with the takeover while it also protects the databases. It acquires the key shared resources by reserving the DASDs that contain the current OLDS and WADS.

After it reserves the log, the alternate IMS performs three activities in parallel:

- Recovers the data in the message queues and the databases
- Performs network changes
- Executes new and reprocessed transactions

The following topics describe these three parallel activities:

- “Recovering data in message queues and databases”
- “Initiating network changes” on page 616
- “Executing new and reprocessed transactions” on page 617

## **Recovering data in message queues and databases**

To process the messages that the active IMS did not begin to process, the alternate IMS ensures that its own copy of the message queue is complete. It can then process these messages and update the databases.

To recover the in-flight transactions, the alternate IMS backs out of the databases all the transactions that the active IMS began but did not complete, and reprocesses the transactions. To ensure integrity of the Fast Path databases, it performs forward recoveries.

The alternate IMS performs I/O toleration by delaying the writing to the databases of the changes to the data block or VSAM control interval that could be overwritten by the failing active IMS until the alternate IMS is certain that the failing active IMS can no longer write to the databases. The alternate IMS keeps the changes temporarily in storage and makes all references to these changes to this storage area rather than to the database. The operators see the term I/O TOLERATION in status displays under the heading MODE. If AVM is not active on the failing active IMS, the operator at the active console must prevent I/O operations by canceling IMS or resetting the system.

The alternate IMS system performs the following actions as part of I/O toleration:

- It identifies and tags database blocks or control intervals that were being updated at the time of failure with an Extended Error Queue Element (EEQE).
- It keeps these records in virtual buffers to enable access during dynamic backout and transaction processing until the user terminates I/O toleration. It also propagates the EEQEs to DBRC and notifies sharing subsystems.
- It backs out possibly incomplete updates that are in-flight and reschedules them for processing.
- It stops using any databases that are being extended at takeover.

IMS restricts access to any data that it shares (through the service of IRLM) with an IMS system outside the XRF complex. Specifically, it prevents other IMS systems from accessing the blocks of data that are affected by in-flight transactions.

When an operator at the alternate IMS replies "GO" to message AVM005A: REPLY "GO" WHEN I/O PREVENTION COMPLETES (or issues the /UNLOCK SYSTEM command):

- IMS updates the databases with the changes for in-flight data it kept in storage.
- IMS makes all portions of the databases available to the IMS systems outside the XRF complex.
- I/O toleration ends.

The IMS command /UNLOCK SYSTEM performs the same action as the response "GO".

## **Initiating network changes**

The changes that occur to the network during a takeover differ depending on whether your XRF complex uses MNPS or USERVAR.

### **Network changes in an XRF complex that uses MNPS**

In the event of an XRF takeover, the alternate IMS system opens a new MNPS ACB. Because the new MNPS ACB uses the same name as the MNPS ACB of the failed IMS system, the MNPS ACB of the failed IMS system must close before the new MNPS ACB can be started. Until the new MNPS ACB opens, VTAM maintains all class-1 terminal sessions open as persistent sessions. After the new MNPS ACB opens, VTAM routes all terminal sessions through the new MNPS ACB.



At takeover, terminal sessions are handled in one of three ways: sessions on class-1 terminals switch to the new MNPS ACB, sessions on class-2 terminals are reestablished after closing briefly, and sessions on class-3 terminals are terminated. For class-3, service is almost the same as if XRF were not present. “Terminals in an XRF complex” on page 633 describes the different characteristics of class-1, class-2, and class-3 terminals.

How quickly a session switches to the alternate IMS or is reestablished depends on the type of CPC in the alternate IMS, the switching priority of the terminal, and the size of the network.

IMS tries to reestablish sessions with the class-2 terminals, such as terminals that communicate with XRF IMS through MSC physical links.

### **Network changes in an XRF complex that uses USERVAR**

To allow additional terminal users to log on to the alternate IMS after an XRF takeover, VTAM must change the application program name in its USERVAR table. The alternate IMS issues a MODIFY USERVAR command to change the entry in its own USERVAR table. Network operators at all other VTAMs that communicate with XRF IMS must issue the MODIFY command to change the application program name in their USERVAR tables if this procedure has not been automated. This procedure can be automated through Tivoli NetView for z/OS. When the alternate IMS issues the MODIFY command, Tivoli NetView for z/OS signals the change in application program name to another Tivoli NetView for z/OS with which it is in session.

At takeover, three obvious changes occur on the terminals: sessions on class-1 terminals switch to backup sessions, sessions on class-2 terminals are reestablished, and sessions on class-3 terminals are terminated. For class-3, service is almost the same as if XRF were not present. “Terminals in an XRF complex” on page 633 describes the different characteristics of class-1, class-2, and class-3 terminals.

The alternate IMS requests that NCP switch each primary session on a class-1 terminal to the preopened backup session. NCP cannot switch a session that does not have a backup session already open. If the failure in the active IMS has not already caused the session with the failing IMS to terminate, NCP then terminates the primary session.

How quickly a session switches to backup or is reestablished depends on the type of CPC in the alternate IMS, the switching priority of the terminal, the size of the network, and the number of NCPs performing the switches.

IMS tries to reestablish sessions with the class-2 terminals, such as terminals that communicate with XRF IMS through MSC physical links.

### **Executing new and reprocessed transactions**

The alternate IMS isolates the log from the active IMS and deletes uncommitted messages in the message queues. It then reacquires the locks needed to perform recovery and starts processing new and reprocessed transactions. As IMS recovers sessions on terminals and makes shared databases available, processing of new transactions increases.

At a takeover, nonshared databases are available immediately. Shared full-function databases become available after completion of the Database Recovery Control (DBRC) reverify. Shared data entry databases (DEDBs) become available after



completion of forward recovery and DBRC reverify, and after a simple checkpoint is taken.

## **Takeover processes without I/O prevention**

The topics “Recovering data in message queues and databases” on page 615, “Initiating network changes” on page 616, “Executing new and reprocessed transactions” on page 617 describe the takeover process caused by the IMS abend and the /SWITCH SYSTEM ACTIVE command. For all other causes of a takeover, AVM does not perform I/O prevention. The IMS operator at the active IMS must perform one of two actions before notifying the IMS operator at the alternate IMS that the active IMS can no longer write to the database.

- If the CPC or z/OS operating system fails, the operator must reset the CPC.
- For all other causes of failure, the operator must terminate IMS to invoke I/O prevention.

Because there is no communication between the active z/OS operating system and the alternate z/OS operating system, an operator at the active z/OS operating system must inform an operator at the alternate z/OS operating system that the databases are safe.

## **IRLM processing at takeover**

The following figure shows an XRF complex that shares its databases with another XRF complex. The primary and secondary relationships refer to the IMS partners. Communication exists between all IRLMs in the sharing group. One complex consists of:

- An active IMS, IMS1, operating with IRLM1
- An alternate IMS, IMS2, operating with IRLM2

The second complex consists of:

- An active IMS, IMS3, operating with IRLM3
- An alternate IMS, IMS4, operating with IRLM4




Figure 70. Two data-sharing XRF complexes

When the operators start these complexes, the IMS systems initialize these sharing relationships:

- Primary sharing partners IMS1 and IMS3
- Secondary partners between:
  - IMS3 and IMS2
  - IMS1 and IMS4
  - IMS2 and IMS4

At a takeover caused by a failure of IMS1, IMS2 notifies IRLM2 of the takeover and tells the operator: IRLM TAKEOVER ISSUED. IRLM2 tells IRLM3 to prepare for takeover and IMS2 acquires the locks that were held by IMS1. When IMS2 completes its takeover processing, it assumes the identity of IMS1. IMS3 and IMS2 (alias IMS1) now become the primary sharing partners. IMS2 issues message DFS3887.

Once the takeover is complete, IMS1 can restart as the alternate IMS on IRLM1.

Data sharing at the block level adds to the post-takeover procedures for the IRLM operator. Before bringing up IMS1 as the alternate IMS, the operator must restart IRLM1 to establish two new secondary IRLM sessions between:

- IRLM3 and IRLM1
- IRLM4 and IRLM1

**Related reading** Similar procedures follow a takeover in the other complex. See *IMS Version 11 Operations and Automation* for descriptions of those procedures.

XRF IMS might use IRLM as a local lock manager. In this case, the IRLMs do not communicate with each other. At takeover, the alternate IMS acquires its own locks and issues takeover messages.

IRLMs can be defined using only the APPL2 and APPL3 parameters. A secondary session is promoted to primary when two active IMS systems identify to two different IRLMs having a secondary session. The promoted session remains primary until a takeover forces reconfiguration or until a valid data sharing situation arises with a third IRLM. If a primary session is disrupted through session failure or loss of one of the IRLMs, that session can be reestablished and automatically becomes primary if another primary session has not already developed.

The two sides of a session must be defined equally, and each session must be unique; that is, two APPL2 sessions cannot be defined. The APPLS parameter overrides the automatic session determination until a takeover occurs.

When a primary session is demoted to a secondary session after a takeover occurs, no operator intervention is required.

## **Practical uses of the planned takeover**

Through the planned takeover, XRF gives you flexibility to perform some tasks during the day that in the past you had to do in evening hours or on weekends. A planned takeover begins when the operator at the active IMS issues a /SWITCH SYSTEM ACTIVE command and the workload on the active IMS transfers to the alternate IMS. The active IMS attempts to quiesce before the alternate IMS performs the takeover. Dependent regions are allowed to reach synchronization points. After the operator brings down the former active IMS, the system programmer can perform these tasks on the former active IMS:

- Testing of backup procedures
- Code maintenance
- Preventive maintenance
- Hardware maintenance
- Library maintenance
- Hardware configuration changes

After you have performed the task that you planned, bring up that system as the alternate IMS in the XRF complex.

You must use the planned takeover carefully; some hardware and IMS initializations must occur simultaneously in both environments. For example, because the active IMS and the alternate IMS must have the same system definition, you cannot request a takeover to change the IMS system definition while the system is the alternate. Unless you can make the system definition changes using online change, a graceful shutdown of both subsystems, followed by an IMS cold start, is required.

In general, you can apply most program temporary fixes (PTFs) without disrupting your end users. Exceptions are when you must terminate the XRF complex to apply the service in both systems simultaneously. These cases are identified in the

PTFs. For IMS service, if the maintenance does not require an IMS cold start in a non-XRF environment or change the IMS log record format, you can apply the PTF while one of the systems is running.

Use the planned takeover to shift the workload to another environment when demands for IMS services increase or decrease. For example, you can run your production work on one CPC during the daytime period of peak activity. You can shift the workload to a different CPC during the nighttime period of lower activity. These two takeovers each day might be a valuable use of XRF for your installation.

## Post-takeover phase of the XRF process

The post-takeover phase overlaps the completion of the takeover phase. As the new active IMS processes new user transactions, IMS continues to perform DL/I back outs, Fast Path forward recoveries, and network switching. The operator can bring up a new alternate IMS after the new active IMS is told that I/O prevention is complete and the old active IMS has terminated.

The following events occur in parallel on the new active subsystem:

- The standard Fast Path DEDB forward recovery is completed. IMS reacquires locks for VSAM CIs containing unwritten but committed changes. An EEQE is created for each of these changes.
- DL/I restart database back outs are completed. The alternate IMS inherits the DBRC authorization for the databases from the previous active IMS instead of being closed during the restart.
- The new active IMS page fixes the non-Fast Path options identified by the DEFERFIX=xx parameter in DFSHSBxx.
- The new active IMS notifies the z/OS availability manager that it is the active IMS. The operator of the failed active IMS must inform the operator of the alternate IMS when I/O prevention is complete so that I/O toleration on the alternate IMS can be discontinued.
- Service is returned to some class-3 terminals.

When the takeover is completed, the alternate IMS becomes the new active IMS, as shown in the following figure.




Figure 71. XRF post-takeover phase

During the post-takeover phase, the new active IMS does not have an alternate environment. It has responsibility for providing service to the end users without benefit of XRF.

The failed IMS system performs problem determination activities. When these are completed, the operator brings up the failed IMS as a new alternate IMS.

If problem determination activity requires lengthy dumping of the IMS data areas, you might consider bringing up a third system as the new alternate IMS in the XRF complex.

The post-takeover phase ends when an operator establishes a new alternate IMS—and thus a new XRF complex.

## Returning the failed active IMS as the new active IMS

Instead of leaving the complex in this arrangement, you might want to return to the original one. For example, if the failed CPC is superior to the other in speed of execution or storage capacity, you might want to restore it as the active IMS as quickly as possible. You cannot re-initialize the failed active IMS as a new active IMS; you must stage its return. First, bring it up as an alternate IMS and then manually request a takeover. This shift of the workload, initiated by the `/SWITCH SYSTEM` command, is another example of a planned takeover.

If XRF IMS shares its databases with other IMS systems, the takeover is a more complicated process that involves the efforts of the IMS operator. Procedures for the operator are described in “IRLM processing at takeover” on page 618.

## Bringing up a third system as an alternate IMS

To limit the time the new active IMS is without an alternate IMS, your operator can bring up a third system as an alternate IMS. You can initialize the new

alternate IMS as soon as I/O prevention is complete and the new active IMS issues message DFS994I to signal that it has taken its first checkpoint. At this time, the operator at the third system can issue the /ERESTART BACKUP command to start an alternate IMS.

You are responsible for ensuring that paths exist to all devices, including DASD and any communication controllers.

## Termination phase of the XRF process

While a takeover is a termination in the sense that it removes the active IMS from service, termination as a separate XRF phase occurs when an operator stops one or both of the IMS systems.

The termination phase, shown in the following figure, returns the XRF complex to two separate and independent environments and stops all activity in the alternate IMS.




Figure 72. XRF termination phase

The three types of termination of the active IMS are:

### Normal Shutdown

In a normal shutdown (/CHECKPOINT FREEZE) of the XRF complex, IMS checkpoint processing occurs. When this has completed, IMS writes a X'06' log record and the alternate IMS terminates automatically. The active IMS tells AVM that processing is terminating with I/O prevention.

### Planned Takeover

A SWITCH SYSTEM command is entered on the active IMS, a quiesce of dependent region processing is done on the active IMS, and a X'06' log record is written indicating that the alternate subsystem is to take over processing. The active IMS communicates to AVM that a normal takeover is occurring.

**Abnormal Termination**

In an abnormal termination of the active subsystem, an attempt is made to purge the log buffer and write a X'06' log record. The active IMS terminates from AVM abnormally.

The two types of termination of the alternate IMS are:

**Normal Shutdown**

The operator can stop the alternate IMS by entering /STOP BACKUP on the alternate IMS or /CHECKPOINT FREEZE or /CHECKPOINT DUMPQ on the active IMS. In both cases, the alternate IMS leaves AVM and specifies I/O prevention.

**Abnormal Termination**

The IMS ESTAE exit is entered and the alternate subsystem terminates abnormally from AVM.

## Cycle of XRF phases

The six phases of XRF form a cycle. When an operator brings up a new alternate IMS, the two systems in the complex have exactly reversed their roles from the initialization phase.

---

## Organization of XRF complexes

The XRF complex can be organized in several ways, each of which has advantages and disadvantages. For example, an XRF complex on one CPC with two IMS systems defined protects IMS from an IMS failure, but does not protect your system from a z/OS or hardware failure.

### One XRF complex with one CPC

An XRF complex can be set up on one CPC with two IMS systems defined. However, the full benefit of XRF IMS is only achieved when IMS as well as z/OS, VTAM, and the CPC are replicated. When only one CPC is used, only IMS has an alternate. Failures in z/OS, VTAM, and the CPC affect both IMS systems.

This configuration consists of:

- One CPC, running z/OS and VTAM
- Two IMS systems, the active and the alternate IMS systems
- IMS system logs shared by the active and the alternate IMS systems
- Databases shared by the active and the alternate IMS systems
- Remote terminals that use SNA protocol and that can communicate with either of the IMS systems
- If your XRF complex uses USERVAR, NCP running in each 37x5 Communication Controller to which the SNA terminals are attached




Figure 73. One XRF complex in one CPC

This arrangement is recommended only as a test environment, and then only if the CPC is fast enough and the system has enough real and virtual storage to support XRF.

## Processing in a one-CPC XRF complex

In an XRF complex that runs on one CPC, XRF processing differs little from XRF running on two CPCs. One important difference is the role of AVM during takeover. In a 2-CPC XRF complex, the alternate IMS issues a RESERVE command to the IMS log to prevent the active IMS from accessing the databases. In a 1-CPC complex, this action does not prevent the active IMS from changing the databases. I/O prevention quiesces the outstanding I/O requests to the system log, as well as to the database data sets. When I/O prevention is complete, the operator can respond “GO” to AVM006A to complete the takeover. If AVM cannot perform I/O prevention, the alternate IMS waits for termination of the failing active IMS before completing the takeover.

## Planning considerations in a one-CPC complex

In a 1-CPC XRF complex, all messages from AVM appear at the one z/OS console as if two z/OS operating systems were in the complex.

If AVM fails to perform I/O prevention at takeover, the IMS operator cannot reset the CPC to prevent the active IMS from writing to the databases. This action terminates the XRF complex. In a 1-CPC complex, the operator should terminate IMS and ensure that all the address spaces in the active IMS terminate before assuming that I/O prevention is complete.



## One XRF complex with two CPCs

An XRF complex in this configuration, illustrated in the following figure, is the most typical.

It requires the following components:

- Two CPCs, operating systems, and VTAMs
- Two IMS systems—one active and one alternate
- Databases shared by the active and the alternate IMS systems
- IMS system logs shared by the active and the alternate IMS systems
- Remote terminals that use SNA protocol and that can communicate with either of the IMS systems
- If your XRF complex uses USERVAR, NCP running in each 37x5 Communication Controller to which the SNA terminals are attached




Figure 74. One XRF complex in two CPCs

## One XRF complex with two CPCs and a non-XRF IMS

An XRF complex in this configuration, illustrated in the following figure, requires both active and alternate IMS systems that share databases.

This configuration, illustrated in the following figure, requires:

- Three CPCs, z/OSs, and VTAMs (only two of the CPCs, z/OSs, and VTAMs are in the XRF)
- Three IMS systems—two active and one alternate
- Databases shared by both active IMS systems and the alternate IMS
- IMS system logs for each active IMS system and one shared system log with the alternate IMS

- Remote terminals using SNA protocol and communicating with one or both of the active IMS systems
- If your XRF complex uses USERVAR, NCP running in the 37x5 Communication Controller that is attached to the CPCs




Figure 75. One XRF complex in two CPCs and a non-XRF IMS

## Two XRF complexes with three CPCs

You can have two XRF complexes running on three CPCs as shown in the following figure, providing adequate real and virtual storage is available on the CPC that runs the two alternate subsystems.

This configuration provides an advantage over two XRF complexes in four CPCs; your alternate IMS systems both run in a single CPC, eliminating the need for one VTAM, one z/OS, and one CPC. This configuration consists of:

- Three CPCs, z/OSs, and VTAMs
- Four IMS systems—two active IMS systems and two alternate IMS systems
- IMS system logs for each active subsystem and shared by its alternate IMS
- Databases shared by the active IMS and alternate IMS in each complex
- Remote terminals using SNA protocol and communicating with one or both of the active IMS systems
- If your XRF complex uses USERVAR, NCP running in each 37x5 Communication Controller that is attached to the CPCs




Figure 76. Two XRF complexes in three CPCs

Paths must exist to both the active and alternate subsystems for users of each subsystem. Each database must have shared access, either online or switchable at takeover, to its associated active and alternate CPC.

## Two XRF complexes with four CPCs

Two XRF complexes that include data sharing must have IRLMs present in all CPCs.

This configuration consists of:

- Four CPCs, z/OSs, VTAMs, and IRLMs
- Two IMS active subsystems, each on a preferred CPC (IMS1 and IMS2)
- Two IMS alternate subsystems, each on a preferred CPC (IMS3 and IMS4)
- Two IMS system logs shared by the active and alternate subsystems
- Databases shared by the active and alternate subsystems
- Remote terminals using SNA protocol and communicating with one or both of the active IMS systems
- If your XRF complex uses USERVAR, NCP running in each 37x5 Communication Controller that is attached to the CPCs

The following figure shows two XRF complexes in four CPCs.




Figure 77. Two XRF complexes in four CPCs

Common access paths are in place on all IMS systems for RECON data sets and all shared databases.

## Planning an XRF complex

You must make some planning decisions about the takeover, such as some of the conditions that force a takeover, the amount of operator involvement, what terminal sessions recover automatically at takeover, and what methods the alternate IMS uses to learn of problems in the active IMS.

## Limitations of XRF

A takeover is not always performed when a failure occurs at an IMS installation.

For example, XRF does not address the outages caused by failures of:

- A channel or link failure that causes a break in communication between the CPC and the network communication controllers
- Failures in the telecommunication network, such as controllers, NCP, lines, and terminals
- Intersystem failures, such as those caused by JES or CTCs
- Some failures in application programs
- Loss of or damage to the IMS databases or log
- A power failure that affects both CPCs in the complex
- Some operator errors

**Related reading:** For information on the z/OS Restart Manager (ARM) in an XRF complex, see “z/OS planning considerations” on page 644.

## Non-XRF workload on the alternate IMS

Processing the takeover significantly increases the work of the alternate IMS. Your installation must prepare to have the processing of the non-XRF workload on the alternate IMS slow down or cease. If the work can be swapped, z/OS might swap it out temporarily. When z/OS swaps the workload back in, it runs in a degraded mode, depending on how demanding the IMS workload is on the CPC.

If you cancel the jobs on the alternate IMS, issue the CANCEL command for one job at a time after the takeover is complete. If you do not wait for the takeover to complete, the takeover processing slows down because CANCEL command processing takes priority over takeover processing.

The non-XRF work that remains on the alternate IMS competes for resources with XRF-related work:

- At takeover, the work that cannot be swapped competes with takeover processing. The system resource manager (SRM) tries to give the alternate IMS the resources to take over the workload from the active IMS. SRM speeds up its analysis of the need for storage on the alternate IMS system. The takeover might cause high paging activity and slower response time for the work that cannot be swapped. However, processing does continue.
- After takeover, the work that cannot be swapped competes with the IMS interactive workload. Response time for the non-XRF users depends on how you code the installation performance specifications (IPS) settings and on the demand for real storage by the work that cannot be swapped.

If you run TSO on the alternate IMS, it should have a lower priority than IMS. At takeover, z/OS swaps out the TSO user address spaces if it needs the real storage. These address spaces remain swapped out until the new active IMS is stable; the duration depends on the installation profile. The alternate IMS should have IPS settings and SRM parameters that are the same as those for the active IMS, in order to favor IMS. When the takeover occurs, the IMS workload is favored.

## Using the Intersystem Communication link

XRF can make effective use of an intersystem communication link (ISC) session between the active IMS and the alternate IMS.

The session is useful because:

- As surveillance, it alerts the alternate IMS to failures in the active IMS.
- During the synchronization phase, it allows the active IMS to take a SNAPQ checkpoint without operator intervention.

The ISC link is a normal VTAM APPL-to-APPL session and can be a route through any of the following:

1. A channel-to-channel protocol (CTC or 3088)
2. A 37x5 Communication Controller with channel paths to both systems
3. Separate 37x5 Communication Controllers through a teleprocessing (TP) line

The following figure illustrates the three kinds of ISC links. The numbers correspond to the numbers in the previous list. Generally, the more hardware units there are on the route, the slower the communication and the greater the chance for failure. Choose the most reliable routes you can for surveillance and other communication.




Figure 78. ISC link options

## Terminals in an XRF complex

XRF supports all of the terminals that presently serve IMS users. It gives this support with a minimum of effort from you. XRF IMS requires that you specify keywords on four or five system definition macro statements or ETO logon descriptors.

Unless you want to change the switching priority of a terminal from the defaults or you want to change the terminal type sessions, you need not change any system definition macros.

XRF IMS gives the terminals at your installation the best service that the terminal's characteristics allow. This topic describes these characteristics.

For the purposes of this explanation, the word “terminal” is used to describe a component of a terminal system, including a programmable controller and its attached operator terminals, printers, and auxiliary control units. It also describes remote subsystems and MSC links.



## Class-1 terminals

For you to take full advantage of XRF, most of the terminals in your complex should be class-1 terminals. Class-1 terminals are those terminals that:

- Use SNA protocol
- Are controlled by a VTAM and NCP that support XRF
- If your XRF complex uses USERVAR, are connected to a 37x5 Communication Controller
- Can be defined on the UNITYPE keyword on the TYPE macro during IMS system definition as one of the following:
  - SLUTYPE1
  - SLUTYPE2
  - SLUTYPEP
  - FINANCE

You define a terminal as class-1 by using the `BACKUP=` keyword. The `BACKUP=` keyword is in the `TERMINAL` macro, other system definition macros, and the ETO logon descriptors.

A class-1 terminal is defined by `BACKUP=(n, YES)`, where *n* is a priority setting between 1 and 7. The priority setting tells IMS to reconnect the terminal and the order in which IMS should do so. The YES parameter tells VTAM to track the data traffic on the terminal so that work on the terminal is uninterrupted and, for XRF complexes that use USERVAR only, create a backup session in the alternate IMS.

When a takeover occurs in an XRF complex that uses MNPS, the alternate IMS system pauses all data traffic on class-1 terminals and maintains persistent sessions for the class-1 terminals until the alternate IMS system opens its own MNPS ACB. The alternate IMS then reconnects the class-1 terminals through the new MNPS ACB. Any data traffic in progress prior to takeover is resumed. Note that the MNPS ACB of the active IMS must be closed prior to opening a new one on the alternate IMS.

In an XRF complex that uses USERVAR, class-1 terminals have backup sessions on the alternate IMS system. These sessions are established during the tracking phase when a session on the active IMS is initiated. The terminals are switched automatically after a takeover without loss of the session.

Setting a high priority for a terminal does not guarantee that a user will receive earlier session recovery at takeover. Because of the competition for resources and transmission speed across lines, recovery time is unpredictable.

Generally, the degree of transparency at takeover depends on the:

- Type of terminal
- Activity on the terminal at a takeover
- Session recovery priority assigned to the terminal

For example, a FINANCE terminal with no output or input in progress at takeover and a `BACKUP=7` priority has the most transparent recovery.

From a user's point of view, a class-1 terminal session that uses MNPS remains unaffected during a takeover; however, from XRF's point of view, at the initiation of a takeover, the class-1 terminal session is connected to a failed IMS. Once

takeover begins, VTAM holds the class-1 terminal session open as a persistent session, the original MNPS ACB closes, and the alternate IMS opens its own MNPS ACB. VTAM then reconnects the class-1 terminal session to that new MNPS ACB. VTAM continuously tracks the session state and data traffic of class-1 terminals. During the takeover process, VTAM pauses any information that was in transit to or from class-1 terminals prior to takeover. When takeover is complete, VTAM resumes the transmission of the information.

From the terminal user's point of view (for terminals that use USERVAR), only one session with IMS exists. From the NCP point of view, two host sessions exist. The actual messages are passed only between the logical unit (LU) and the active IMS. The alternate IMS tracks the status of the message activity by monitoring the IMS log. When the active IMS establishes or terminates a terminal session with class-1 terminals, the alternate IMS establishes or terminates a backup session. When the active IMS fails, the alternate IMS takes over any open terminal sessions without losing control from the viewpoint of the LU by changing the mode of the pre-established backup session from BACKUP to ACTIVE. When the session is switched, the NCP sends the alternate IMS its view of the terminal's status. IMS compares this with its own record of status to decide what, recovery action to take, if any. When the REVERIFY operand is used with RACF, the user must sign on again.

## **Class-2 terminals**

Class-2 terminals do not have backup session support on the alternate IMS, but are restarted automatically after takeover. The interface to the terminal, including recovery procedures after session reestablishment on the alternate IMS, are the same whether your XRF complex uses MNPS or USERVAR.

You define a terminal as class-2 using the BACKUP= keyword. The BACKUP= keyword is in the TERMINAL macro, other system definition macros, and the ETO logon descriptors.

A class-2 terminal is defined by BACKUP=(*n*, NO), where *n* is a priority setting between 1 and 7. The priority setting tells IMS to reconnect the terminal after takeover and the order in which IMS should do so. The NO parameter specifies that VTAM will not track the data traffic on the terminal.

When a class-2 terminal session is established on the active IMS, the alternate IMS tracks the session initiation or termination using the log records. When the active IMS terminates abnormally, the alternate IMS tries to establish a new session with the network resources that were active before the failure.

Class-2 terminals receive good support from XRF. In fact, at takeover, IMS might be able to reestablish service on some of your class-2 terminals before all of your class-1 terminal sessions are switched.

When IMS tries to reestablish a session with a terminal to which there is no available path, the attempt fails. Before IMS can reestablish a session on a locally attached terminal, the operator must switch the terminals to the new active IMS. If the complex has many of these kinds of terminals, you might consider giving the operator control at takeover before the takeover actually proceeds. At this time, the operator can perform the switch, allowing a faster recovery for sessions on these terminals. You give the operator this control in the AUTO parameter in member DFSHSBxx in IMS.PROCLIB.

Although IMS tries to establish new sessions for the class-2 terminals, if the owner of the terminals or the link to the owner fails, the attempt to establish the new session also fails.

Terminals that qualify as class-2 are:

- Multiple Systems Coupling (MSC) and ISC subsystems that communicate with IMS XRF through VTAM
- Spool line groups on shared DASD, locally attached to both CPCs
- Non-SNA 3270 terminals controlled by VTAM
- Devices controlled by the Network Terminal Option (NTO) licensed program
- Locally attached devices
- Terminals that do not use the SNA protocol
- Terminals that qualify as class-1 terminals, except:
  - The terminals are not controlled by a 37x5 Communication Controller
  - The terminals are not controlled by an NCP or VTAM that supports XRF or defined BACKUP=(n,NO) on the system definition macro or ETO logon descriptor

If you have an application terminal that communicates with a programmable controller in an XRF complex, you might need to add code to the communications portion to achieve the highest level of transparency available. Additionally, for the ISC terminal, you must define a session for the link that connects the ISC terminal to XRF IMS.

**Related reading** For information on what code you need to add to the programs, see *IMS Version 11 Communications and Connections*.

### Class-3 terminals

Class-3 terminals do not have backup session support on the alternate IMS. Their terminal sessions must be restarted manually on the new active IMS after takeover, either by the MTO or by user logon.

Terminals that are eligible for class-3 include all the terminals eligible for class-3 service (such as all LU 6.2 devices), and class-1 and class-2 terminals that have indicated no switching should be done. This is done by specifying BACKUP=NO on the system definition macro or ETO logon descriptor.

Class-3 terminals must be physically connected to both the active and alternate IMS systems. If they are not physically connected to both IMS systems, reestablishment of sessions might not be possible.

### Specifying terminal support

All terminals that IMS supports can be used in an XRF complex. The terminal characteristics determine the type of support a terminal can have in an XRF complex, with class-1 having the best support and class-3 having no support.

LU 6.2 devices are always class-3. ETO terminals and users are eligible for any level of support.

**Related reading** For more information on LU 6.2 and ETO in XRF, see “Specifying terminal support” in *IMS Version 11 Communications and Connections*.

The following table summarizes the eligibility of various devices for the terminal classes.

*Table 47. Class eligibility for various devices*

| Device                            | Eligible for class-1 | Eligible for class-2 | Eligible for class-3 |
|-----------------------------------|----------------------|----------------------|----------------------|
| SLUTYPE1 (3286,...)               | Yes                  | Yes                  | Yes                  |
| SLUTYPE2 (3278,3279,...)          | Yes                  | Yes                  | Yes                  |
| SLUTYPEP (4700,8100,SERIES 1,...) | Yes                  | Yes                  | Yes                  |
| FINANCE (3601)                    | Yes                  | Yes                  | Yes                  |
| LUTYPE6 (ISC)                     | No                   | Yes                  | Yes                  |
| MSC (VTAM)                        | No                   | Yes                  | Yes                  |
| MSC (BSC)                         | No                   | Yes                  | Yes                  |
| NTO (leased line)                 | No                   | Yes                  | Yes                  |
| 3270 (VTAM)                       | No                   | Yes                  | Yes                  |
| Spool line groups                 | No                   | Yes                  | Yes                  |
| Locally attached devices          | No                   | Yes                  | Yes                  |

To change the priority or terminal type of the devices listed in the table, use one of the following BACKUP keywords on the TERMINAL Macro or ETO logon descriptors:

- For devices eligible for class-1, class-2, or class-3 service, use:
  - BACKUP=(1-7, YES): for class-1 service
  - BACKUP=(1-7, NO): for class-2 service
  - BACKUP=NO: for class-3 service
- For devices eligible for only class-2 or 3 service, use:
  - BACKUP=(1-7, NO): for class-2 service
  - BACKUP=NO: for class-3 service

The BACKUP= keyword on several macros in the system definition statement specifies the support a terminal receives. Your installation determines which terminals have backup sessions established and the switching priority of the terminals. Level 1 is the lowest priority, and level 7 is the highest. Although VTAM requests are priority ordered by IMS, the requests can be completed in any order because of unpredictable factors in the network, such as pacing, line speeds, and congestion.

Terminals defined with BACKUP=(1-7,YES) are automatically taken over without losing session control or any data inflight at the time of failure. This constitutes class-1 support.

In XRF complexes that use USERVAR, terminal backup sessions are established during the tracking phase for terminals defined with BACKUP=(1-7,YES). You specify the maximum number of terminal backup sessions that can be established on the alternate IMS system in the NCP BUILD statement by using the BACKUP=n option, where n is the maximum number of terminal backup sessions. When IMS reaches the maximum number of terminal backup sessions, IMS denies logon to

any additional class-1 terminal users. Carefully determine the maximum number of backup sessions that should be on the alternate IMS, and specify an appropriate number.

If BACKUP=(1-7,NO), this terminal is automatically restarted after takeover, but no BACKUP session is established. This forces class-2 support.

For both class-1 and class-2 terminals, the default switching priority is 4.

If BACKUP=NO, this terminal must be restarted manually after takeover. This forces class-3 support.

## How takeover affects a terminal user

A user's awareness of a takeover depends on the class of terminal, the priority with which IMS requests session recovery, the access method that controls the terminal, and the status of the input or output message at the time of failure. The main network objective is to recover service on the terminals as soon as possible.

As sessions on class-1 and class-2 terminals recover, users might notice a pause in operations. If a user on a class-1 terminal has just entered a message, the alternate IMS might request that the user reenter the message or issue message DFS386II. In both cases, the user should clear the screen and reenter the last message.

**Related reading** For more information on how to ensure that users of FINANCE and SLUTYPEP terminals receive this message, see *IMS Version 11 Communications and Connections*.

IMS notifies most class-2 terminal users of the recovered sessions by issuing message DFS2469W, DFS3649, or DFS3650, indicating that the session has been reconnected. FINANCE, SLUTYPEP, and ISC terminals do not receive these messages.

**Related reading** For more information on these messages, see *IMS Version 11 Messages and Codes, Volume 1: DFS Messages*.

The IMS operator must manually recover service on terminals that fail at a takeover.

### Takeover for class-1 terminal users (except SLUTYPE2 users)

If the user is between transactions, has received a reply to the last data entry, and has not yet begun the next transaction, no indication exists that a takeover has occurred and that the terminal has been switched. The user might notice that a takeover has occurred if a failure occurs during input, output, or when transactions are in-flight.

#### Input

##### Queued Recoverable

If the input message of a recoverable transaction has been queued, it is automatically scheduled after takeover and the user is not aware of the terminal switch.

##### Not Yet Queued Recoverable

If the input message of a recoverable transaction had not been queued at the time of the failure, the message is lost. This includes Fast Path transactions that had not reached a sync point.

**Conversational**

If the transaction was conversational, the alternate IMS sends an Exception Response (SSENSE=X'0826', USENSE=X'0F15') message. DFS3861I is also sent if the terminal can receive system messages. The RTO can issue /HOLD and /RELEASE commands for the conversation and receive the first page of the final output again.

**Non-conversational**

If the transaction was non-conversational and the last output message was not MFS, the Exception Response and system message are sent (see Conversational, above), and the RTO resumes with a new transaction.

**Queued Nonrecoverable**

If the transaction is nonrecoverable and has been queued, the message is lost.

**Not Yet Queued Nonrecoverable**

If the transaction is nonrecoverable and has not been queued, the message is lost. The alternate IMS sends an Exception Response (SSENSE=X'0826', USENSE=X'0F15') and message DFS3861I.

**In-Flight****Recoverable**

If the transaction is being processed at the time of failure, IMS backs out the updates, reschedules the transaction, and sends the appropriate output message after completion of the transaction. The takeover is transparent to the user only if there is no delay in scheduling the transaction. Otherwise, the terminal user experiences a response-time delay. The user does not need to reenter any data.

**Nonrecoverable**

If the transaction is nonrecoverable or is a Fast Path transaction that has not reached sync point, the transaction is lost. The Exception Response (SSENSE=X'0826', USENSE=X'0F15') and message DFS3861I are sent.

**Output****Recoverable**

If the output message at the time of failure is recoverable, the alternate IMS continues to send the message from the next segment, and the terminal user is not aware of the takeover.

**Nonrecoverable**

If the output message at the time of failure is nonrecoverable, the alternate IMS resets the current message and sends message DFS3861I if the terminal can receive system messages.

**Unknown Output**

If the alternate IMS does not know the contents of the last output message (if the last output was a system message which is not logged), the alternate IMS cancels the current output message and sends message DFS3861I if the terminal can receive system messages.

## Takeover for class-1 SLUTYPE2 terminal users

Terminal users at SLUTYPE2 devices are always aware of takeover activity. If the terminals are defined to receive system messages, they receive message DFS3861I text after the takeover. I/O activity at the time of failure is always reset.

### Input

#### Queued Recoverable

If the input message of a recoverable transaction is queued, it is automatically scheduled after takeover. The terminal user must press the PA1 or PA2 key to receive the output if the terminal is defined to receive DFS3861I messages.

#### Not Queued Recoverable

If the input message of a recoverable transaction is not queued at the time of the failure, the message is lost. This includes Fast Path transactions that had not reached a sync point.

##### 1. Conversational

If the transaction is conversational, the terminal user can receive the first page of the final output by entering /HOLD and /RELEASE commands.

##### 2. Non-conversational

If the transaction is non-conversational, the terminal user must begin a new transaction. A /FORMAT command might be needed.

#### Nonrecoverable

If any nonrecoverable transaction or the resultant reply is in the failed system at the time of failure, the transaction or message is lost.

### In-Flight

#### Recoverable

If the transaction is being processed at the time of failure, IMS backs out the updates and reschedules the transaction. The terminal user must press the PA1 or PA2 key to receive the output message. The user might experience a response-time delay if the transaction is not able to be scheduled immediately after takeover. The user does not need to reenter the transaction.

#### Nonrecoverable

If the transaction is nonrecoverable or is a Fast Path transaction that has not reached sync point, the transaction is lost.

### Output

#### MFS Logical Paging

If the terminal user is paging through logical pages, the user must press the PA1 or PA2 key to get the first screen of the final set of logical pages that are being reviewed.

#### Single MFS Page or Non-MFS Output—Not Dequeued

If the failure occurred between the time IMS sent the output and the time IMS received an acknowledgement of receipt (the message is not dequeued), the final output is sent to the terminal again.

#### Single MFS Page or Non-MFS Output—Dequeued

If the failure occurs after the output is sent and the



acknowledgement is received, the user must begin the transaction again. A /FORMAT command might be needed.

**MFS Multipage Output without Operator Logical Paging—Last Page Sent and Not Dequeued**

If the failure occurs between the time IMS sent the output and the time IMS received an acknowledgement (the message is not dequeued), the first page of the final output is sent again.

**Last Page Sent—Dequeued**

If the failure occurs after the output is sent and the acknowledgement is received, the user must begin the transaction again. A /FORMAT command might be needed.

**Last Page Not Sent**

If the last page is not sent at the time of failure, IMS sends the first page of the output again. The terminal user pages with the PA1 key.

## **Takeover for class-2 terminal users**

Class-2 terminal users have the same support that they have in a non-XRF environment after an emergency restart. The difference is that the LOGON sequence has been handled automatically by takeover. Users at FINANCE, SLU P, and ISC terminals are signed on automatically at takeover. Other ETO users might be required to sign on again after takeover.

**3270 VTAM**

After takeover, IMS attempts to establish a new session on the alternate IMS.

**LUTYPE6 (ISC)**

After takeover, IMS attempts to establish a new ISC session on the alternate IMS.

**MSC (VTAM)**

After takeover, IMS attempts to reestablish all VTAM physical links that were active at the time of failure.

**NTO** After takeover, IMS attempts to establish a new session on the alternate IMS.

**Locally attached devices**

These devices must be switched manually.

**LU 6.2 devices**

Conversations in process must be reallocated.

**Remote devices**

VTAM terminals that are twin tailed or that are local-to-local NCP connected are switched automatically. If the terminal is not physically connected to the new active IMS, the switch fails.

If the owning VTAM fails, the operator must switch the terminals to the new active IMS. Device switching can also be done using a 2914/3814.

**Spool line groups**

After takeover, IMS opens a new spool data set.



## How VTAM ownership affects terminal switching

Terminals can be owned by a VTAM residing in the active CPC, in the alternate CPC, or in a CPC of a communication management configuration (CMC). Having your terminals owned by a host outside the XRF complex offers the most stable ownership, because the host is not affected by takeover, and the takeover is not affected by the CMC activities.

The following table summarizes the terminal's session status, logon capabilities, and logoff capabilities after a takeover. It assumes that VTAM in the active IMS is no longer functioning.

*Table 48. Terminal capabilities of owning VTAM during and after a takeover.*

| System                      | Active |     |     | Alternate |     |     | CMC or other host |     |     |
|-----------------------------|--------|-----|-----|-----------|-----|-----|-------------------|-----|-----|
| Terminal class              | 1      | 2   | 3   | 1         | 2   | 3   | 1                 | 2   | 3   |
| Session continue            | Yes    | No  | No  | Yes       | No  | No  | Yes               | No  | No  |
| LOGOFF enabled              | Yes    | N/A | No  | Yes       | N/A | No  | Yes               | N/A | N/A |
| LOGOFF enabled <sup>1</sup> | Yes    | Yes | Yes | Yes       | Yes | Yes | N/A               | N/A | N/A |
| LOGON enabled               | No     | No  | No  | Yes       | Yes | Yes | Yes               | Yes | Yes |
| LOGON enabled <sup>1</sup>  | Yes    | Yes | Yes | Yes       | Yes | Yes | N/A               | N/A | N/A |

**Note:**

1. After another system acquires terminal ownership

## Planning for VTAM ownership of class-1 terminals

When determining where to place the ownership of the class-1 terminals, consider what causes the least operator effort, the least disruption to class-1 terminal users, and the least strain on the CPC.

Before you make the decision, consider some of the implications of VTAM failures on ownership. First, review some information about VTAM that has special meaning for XRF.

Ownership is significant when a terminal user logs on to IMS or logs off from IMS. Each VTAM has one system services control point (SSCP); it is responsible for session initiation and session termination for terminals that it owns. Therefore, if the owning VTAM fails, some functions of VTAM are no longer available. Service on the class-1 terminals that are logged on to IMS does continue. But when a user of a class-1 terminal within its domain tries to log on, VTAM cannot honor the request. Upon restarting VTAM, the network operator must also reestablish the ownership of the terminals. Alternatively, the network operator can transfer ownership to another VTAM.

At a takeover, the ownership of class-1 terminals does not shift to the VTAM of the alternate IMS system. If you assign ownership to the active VTAM and the active IMS system fails, the class-1 terminals switch to the alternate IMS system, but the VTAM of the failed IMS system retains ownership. The VTAM of the failed IMS system retains ownership even if the failed IMS system returns to service, but functions only as the alternate IMS system.

In making the decision about terminal ownership, consider three options. The owner of the class-1 terminals can be:

- The VTAM in the active system

- The VTAM in the alternate system
- A VTAM in a Communication Management Configuration (CMC)

You should assign the ownership of the class-1 terminals to a third VTAM in a CMC. Having the owning VTAM in a CPC that is outside the XRF complex offers the following advantages:

- It reduces contention on the active system during normal operations.
- It reduces contention on the alternate system during the takeover.
- It prevents failures in the XRF complex from terminating the VTAM that owns the class-1 terminals.
- It allows terminals that are connected to the NCP on SNA switched (that is, dial-up) lines to be class-1 terminals.

If you do not have a CMC, assign the ownership of the class-1 terminals to either VTAM in the complex, and do not change the ownership after a takeover unless you are restarting VTAM.

**Related concepts:**

 z/OS: SNA operations introduction

## NCP planning considerations for XRF with USERVAR

The primary planning task for Network Control Program (NCP) is deciding how many class-1 terminals you want to attach to each 37x5 Communication Controller. For each NCP, specify this number on the BACKUP operand on the BUILD definition statement. This action allows the NCP, through the SSP, to generate the appropriate number of control blocks for primary and backup sessions for class-1 terminals.

The 37x5 Communication Controllers connected to your class-1 terminals need not be dedicated to the IMS work. They can also support other network traffic, including gateway. As you plan the storage requirements for the gateway 37x5 Communication Controller that supports class-1 terminals, remember that this type of 37x5 Communication Controller requires more storage than a 37x5 Communication Controller without XRF terminals and without gateway.

**Related reading:** For more information on NCP storage, see *z/OS Communications Server: SNA Network Implementation Guide*.

## Performance considerations

When you evaluate a configuration with XRF capability, you should distinguish between general availability criteria and response-time objectives. By carefully planning the configuration, you can provide alternative access paths and reduce the single point-of-failure occurrences.

If you consider the total system response characteristics, you evaluate the response for each system when active. Your initial parameters for IMS, page fixing, and region utilization are sized to the CPC and the expected workload.

### Response-time objectives

To account for special situations, you might need to adjust the objectives you have established for the workload on either CPC. For example, exclude from your statistics the situation in which a successful takeover occurs from a planned outage. If you set any objectives for critical

transactions, remember that takeover processing is normally followed by a period when response is stabilizing. An end user might experience a slight delay.

### **Capacity planning**

When planning the virtual storage requirements in an XRF complex, the key question is: how much resource is required to give acceptable support to a given IMS user set? Careful estimates should be made if more than one IMS alternate subsystem is to be run on the same alternate CPC.

---

## **z/OS planning considerations**

z/OS systems without XRF can coexist with z/OS XRF systems in a single installation. In particular, z/OS XRF and non-XRF systems can participate in a single global resource serialization (GRS) ring and in a single JES2 or JES3 complex. Even in one z/OS system, an XRF IMS system can coexist with non-XRF jobs or subsystems, including another IMS.

The following list describes some of these z/OS planning considerations.

### **z/OS Automatic Restart Manager in an XRF complex**

In an XRF environment, when the alternate IMS system has started the tracking phase, it issues a z/OS Automatic Restart Manager (ARM) associate. This ensures that the active IMS system is not automatically restarted by ARM after the backup takes over. If the active IMS were to automatically restart after the backup takes over, the IMS log (OLDS) and the message queue might be destroyed.

If IMS abends before it completes restart, the abended IMS de-registers from ARM and is not automatically restarted, because XRF tracking is considered to have completed restart for an XRF backup.

### **Global resource serialization considerations**

XRF does not require global resource serialization. However, if one CPC in an XRF complex is in a global resource serialization ring, the other CPC in the complex should be in the same ring. z/OS with XRF and z/OS without XRF can participate in a single global resource serialization ring.

XRF and the XRF takeover do not affect the normal procedures for global resource serialization. However, if your two systems are in a global resource serialization ring, you should be aware of the relationship between XRF and global resource serialization.

If the CPC or z/OS operating system the active IMS system is running fails, the alternate IMS system requests a takeover. Messages from global resource serialization indicate that the ring is disrupted. When the system restarts (either automatically through the RESTART(YES) parameter in the GRSCNFxx member of SYS1.PARMLIB or through the operator issuing the VARY GRS,RESTART command), global resource serialization rebuilds the GRS ring that was disrupted. To release the resources owned by the failed system, the operator issues the VARY GRS,PURGE command. When the operator IPLs z/OS and restarts IMS as the new alternate subsystem for the XRF complex, the system rejoins the ring.

If the CPC or z/OS fails in the alternate system, processing on the active IMS continues normally. The operator should enter the VARY GRS,PURGE command to free the resources owned by the failed active IMS and relocate the alternate IMS on another CPC in the ring.

Include all IMS data set names in the global resource serialization SYSTEMS exclusion resource name lists (RNLs). Depending on the naming conventions at your installation, you might be able to include them with one generic entry. If you do not list the names in the SYSTEMS exclusion RNLs, global resource serialization serializes access to these data sets, a service that DBRC and the IRLM already provide the IMS data sets.

Do not include the DBRC RECON or the OLDS and WADS names in the RESERVE conversion RNL. Doing so results in physical reserves being sent to DASD when reserves are issued by IMS.

Backup CTCs are recommended, because they allow an installation to continue using GRS when a CTC fails. Without a backup CTC, all global ENQ requests remain suspended until the operator reestablishes communication between systems.

#### **JES considerations**

z/OS with XRF and z/OS without XRF can both participate in a JES2 or JES3 complex. The active and alternate IMS systems can be in two JES3 global systems. If you use the services of JES3, include all IMS data sets and databases in the RESDSN statement. If you do not do so, and the JES3 global fails, the active IMS cannot allocate the data sets and databases until a JES3 global becomes available.

The IMS data sets must be available when JES3 is brought up.

#### **RACF considerations**

Because the alternate IMS system needs the security information in the RACF data sets at takeover, place the RACF data sets on DASD shared by the active and alternate IMS systems. To avoid single points of failure, use the RACF backup facility to keep a second copy of these data sets also on shared DASD.

The REVERIFY operand on the RACF RDEFINE APPLDATA command is nullified during a takeover. If your installation uses this operand, the users must sign on again after a takeover to identify the password to the new active IMS.

---

## **Preparing the system for XRF**

You define IMS to be part of an XRF complex during the IMS system definition.

Do this by:

- Tailoring the IMS execution JCL
- Coding system definition macro statements to:
  - Establish XRF to be part of the generated system
  - Specify the VTAM application names and passwords
  - Define the IMS master terminal and secondary terminals
  - Customize individual terminals
  - Optionally declare priorities that affect the order in which terminals and MSC links are made available
- Using IMS.PROCLIB members to initialize system data sets

## Tailoring the IMS execution JCL

Each CPC in an XRF complex has an IMS job running continuously. z/OS START commands use two different procedures to initialize the active IMS and the alternate IMS.

The XRF-related parameters in the control region EXEC statement are as follows:

### **HSBID=**

Specified as HSBID=1 in one procedure and HSBID=2 in the other. HSBID=null deactivates XRF capability. If you bring up a third IMS as an alternate IMS after a takeover, specify HSBID= the same as on the failed active IMS.

### **HSBMBR=**

Specified as a 2-character suffix that corresponds to a DFSHSBxx member included in IMS.PROCLIB. This member is where the control information and takeover criteria are specified.

### **MNPS=**

Defines XRF as using MNPS and specifies the name of the MNPS ACB. The MNPS= keyword in the EXEC statement overrides the MNPS= and USERVAR= keywords in the DFSHSBxx PROCLIB member.

### **MNPSPW=**

Defines a password to be used with the MNPS ACB name that is specified in the MNPS= keyword. This is an optional keyword that overrides specifications that are made in the DFSHSBxx PROCLIB member. If VTAM is configured to check passwords, it is this password that VTAM checks. If it is not specified and VTAM expects one, the MNPS ACB is not opened.

### **USERVAR=**

Specifies the USERVAR name. If an MNPS ACB name has been defined using the MNPS= keyword, this keyword is ignored.

## Coding IMS system definition macro statements for XRF

When you define the IMS system, you specify certain keywords on the macro statements that are in your existing generation deck.

XRF keywords appear in the following system definition macros:

COMM  
CTLUNIT  
IMSCtrl  
LINE  
LINEFRP  
MSLINK  
MSPLINK  
STATION  
TERMINAL  
TYPE

For the syntax of the system definition macro statements, see *IMS Version 11 System Definition*.

XRF requires that you code the keywords that:

- Establish XRF support for IMS

- Provide two VTAM application names
- Name two passwords for ACB security
- Define your master and secondary IMS terminals

You can use the defaults on other keywords that define what happens to terminals at a takeover. The defaults allow you to change IMS support at takeover and change the priority that IMS uses to recover sessions at takeover. If you allow XRF to use the defaults on the macros and ETO descriptors that define your terminals, you receive the best recovery IMS can give your terminals. The default priority for IMS to recover sessions at takeover is 4, with 7 being the highest priority.

## Specifying the VTAM application names and the passwords

To specify the VTAM application name and password for the active IMS system and the alternate IMS system, use the APPLID= and PASSWORD= keywords on the COMM macro.

- Specify the VTAM application name for the active IMS system and the alternate IMS system on the APPLID= keyword on the COMM macro.

The APPLID= keyword specifies the application name (APPLID name) by which VTAM knows an IMS system. The APPLID names are stored in ACBs.

APPLID names are used differently depending on whether your XRF complex uses MNPS or USERVAR.

### In an XRF complex that uses MNPS

The APPLID names uniquely define each IMS in the XRF complex to VTAM. The APPLID names can be used to directly logon to a specific IMS system in the XRF complex. The APPLID names and ACBs do not play a role in terminal switching during a takeover in an XRF complex that uses MNPS.

### In an XRF complex that uses USERVAR

- VTAM places the APPLID names in the USERVAR table. Only one APPLID name is in the USERVAR table at any one time: the name that applies to the active IMS system. At takeover, IMS tells VTAM to replace it with the APPLID name of the alternate IMS system.
  - The APPLID names must be the same as the labels on the VTAM APPL definition statement that authorizes VTAM to support backup sessions for class-1 terminals.
- Specify the VTAM password for the active IMS system and the alternate IMS system on the PASSWORD= keyword on the COMM macro.
- The PASSWORD= keyword on the COMM macro defines two passwords that VTAM checks in the ACB.

For example, to specify VTAM application names of IMSNAME1 and IMSNAME2 with passwords of IMSP1 and IMSP2:

```
COMM APPLID=(IMSNAME1,IMSNAME2),PASSWORD=(IMSP1,IMSP2)
```

## Defining the IMS master and secondary terminals

Both the active and alternate IMS systems have IMS primary and secondary master terminals that are always active. To define them, specify the XRF-related keywords on the TERMINAL macro.

If VTAM controls your IMS master terminals, specify the NAME keyword.

The NAME keyword on the TERMINAL macro defines a second node name for a VTAM master or secondary terminal. It is also used to define an IMS-managed ISC link between the active and alternate subsystems when both names match the APPLID keyword on the COMM macro.

### Customizing individual terminals

Regardless of the number of terminals in your complex, obtaining XRF support for those terminals does not require you to specify any additional keywords. By accepting the defaults, IMS determines the best recovery possible for your present terminals.

For example, terminals that meet all of the following criteria are automatically defined as class-1 terminals:

- Defined as SLUTYPE2 on the UNITYPE keyword
- Connected to 37x5 Communication Controllers
- Controlled by NCP and VTAM that support XRF

IMS automatically considers as class-2 any terminal that is defined UNITYPE=SLUTYPE2 and is attached to a 37x5 Communications Controller.

The following table shows the keywords you code to customize individual terminals.

*Table 49. XRF system definition macro keywords for VTAM-controlled terminals*

| Macro    | Keyword  |
|----------|----------|
| COMM     | APPLID   |
|          | PASSWORD |
| TYPE     | BACKUP   |
| TERMINAL | BACKUP   |
|          | NAME     |
| NAME     |          |

The BACKUP keyword on the TYPE macro, TERMINAL macro, and ETO descriptors specifies session recovery on class-1 and class-2 terminals and sets the priority in which IMS recovers service on these terminals. The default for the BACKUP keyword is (4,YES). This means terminals that qualify as class-1 automatically have backup sessions at takeover with a switching priority of 4. Terminals that qualify as class-2 terminals automatically have new sessions established at takeover with a priority of 4.

Specify BACKUP=NO if you have some reason to change class-1 or class-2 support to class-3.

Use the BACKUP keyword to establish a priority other than 4 for recovery of service on your class-1 or class-2 terminals. Specify BACKUP=n, where n is a number between 1 (slowest recovery) and 7 (fastest recovery).

Although these priorities are used by IMS, because of internal VTAM processing, the requests for network connection might be completed in a different order.

To change class-1 support to class-2 support, code BACKUP=n,NO where n is the session recovery priority for the terminal.



For example:

- If you want all the terminals defined by a TYPE macro to be class-2 terminals with the session recovery priority of 7, specify:  
TYPE . . . .BACKUP=(7,NO)
- If you do not want IMS to recover sessions on a class-1 terminal at a takeover, specify:  
TERMINAL . . . .,BACKUP=NO

## Defining MSC links

If an XRF IMS system communicates with another IMS system over MSC links, you can change the priority with which IMS establishes new sessions on these links.

The BACKUP keyword on the MSPLINK macro defines MSC physical link recovery and sets the priority for the establishment of new sessions on these physical links at takeover. Consider raising the priority that IMS uses to establish new sessions across these links. The default is BACKUP=4.

You can use the BACKUP keyword on the MSPLINK macro to override the BACKUP option you specified on the MSPLINK macro for the logical link definition. The default is BACKUP=4.

## Defining IMS.PROCLIB members for XRF

Specify keyword values for XRF in the following PROCLIB data set members: DFSDCxxx, DFSHSBxx, DFSFIXxx, and DFSVSMxx.

The following PROCLIB members include parameters that you can use to define an XRF system:

- DFSDCxxx

The DFSDCxxx PROCLIB member holds the PSTIMER parameter. The PSTIMER parameter specifies the amount of time you want the VTAM to hold open persistent sessions while waiting for the alternate IMS to open a new MNPS ACB. If a new MNPS ACB is not opened before the time runs out, the sessions are closed. The default value is 3600 seconds.

- DFSHSBxx

The DFSHSBxx member specifies the name of the XRF complex as it is known to the z/OS availability manager, to DBRC, and to IRLM. It also specifies the surveillance mechanism and the takeover conditions for the complex.

**Recommendation:** Use the same member for both the active and alternate subsystems.

For details about the parameters for XRF in the DFSHSBxx member, see “XRF parameters in DFSHSBxx” on page 650.

Each subsystem has IMS running continuously. You initially declare which IMS is the active IMS with the HSBID parameter in the EXEC statement. The HSBID parameter, HSBID=1 or HSBID=2, points to the respective IMS APPLID in the COMM statement, the NODE for a VTAM master terminal. The HSBID also identifies all the message queue data sets associated with each subsystem.

You also need to specify, with the HSBMBR parameter, a suffix that corresponds to a DFSHSBxx member included in the IMS.PROCLIB.



## XRF parameters in DFSHSBxx

When you tailor the IMS system for XRF, you must specify an IMS control statement in member DFSHSBxx of IMS.PROCLIB. The control statement consists of a number of parameters.

There is one copy of DFSHSBxx for the active IMS and one for the alternate IMS. The two copies need not be identical, but the parameters RSENAME and either MNPS or USERVAR must be the same in each copy. HSMBR=xx on an IMS.PROCLIB procedure statement defines the DFSHSBxx member.

Although IMS provides commands that operators can use to dynamically change the surveillance options in DFSHSBxx, other parameters in DFSHSBxx change only when z/OS starts a new IMS.

The following operating parameters that you specify in IMS.PROCLIB member DFSHSBxx improve performance:

- SURV
- RSENAME
- MNPS
- MNPSPW
- USERVAR
- SWITCH
- ALARM
- AUTO
- KEYEVENT
- DEFERFIX
- LNK
- LOG
- Restart data set (RDS)

### **SURV=ALL, LNK, LOG, RDS, NO**

XRF IMS provides surveillance to inform the alternate IMS of failures in the active IMS. Without surveillance, some failures in the active IMS occur without any warning to the alternate IMS. The SURV parameter specifies what surveillance mechanisms are to be in effect for the XRF complex.

Options on the SURV parameter are:

#### **Option Result**

- |            |                                                                                                                                                                                                                                                                                                                                                     |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ALL</b> | The ALL parameter indicates that all three surveillance mechanisms (LNK, LOG, and RDS) are to be used.                                                                                                                                                                                                                                              |
| <b>LNK</b> | The LNK parameter indicates that the active IMS is to send signals across an ISC link to the alternate IMS at the frequency specified by the interval value on the LNK parameter. The ISC link is also used to request the SNAPQ checkpoint of the active IMS from the alternate IMS during synchronization.                                        |
| <b>LOG</b> | The LOG parameter indicates that the alternate IMS is to check the OLDS for new records as often as the interval value on the LOG parameter indicates. This check is in addition to the access to the OLDS by the alternate IMS to maintain a system status on the alternate IMS which reflects the processing that is occurring on the active IMS. |

**RDS** The RDS parameter indicates that the active IMS is to place a time stamp in the shared restart data set (RDS) as often as the interval value on the RDS parameter indicates. The alternate IMS checks the RDS to ensure that the time stamps are being written.

**NO** This parameter indicates that no surveillance is active.

Any combination of LNK, LOG, or RDS can be specified but using all three is recommended. You can dynamically change the surveillance and the interval parameters by commands issued from the master terminal.

Each of the three surveillance mechanisms (LNK, LOG, and RDS) is specified with an interval and a timeout value. On the active IMS, the interval value indicates the frequency at which the surveillance mechanism sends a signal or updates a data set. On the alternate IMS, the interval indicates the frequency at which the signals or time stamps are received. If the timeout value is exceeded, the alternate IMS can initiate a takeover. These time values are only approximate; actual values might be slightly longer.

These surveillance parameters determine:

- What surveillance mechanisms operate in the complex
- How often the alternate IMS receives signals from surveillance (LNK, LOG, and RDS)
- Whether the absence of any of the surveillance mechanisms causes a takeover

The alternate IMS uses these three surveillance mechanisms in different ways. The ISC link (LNK) is a low-overhead surveillance mechanism that is best used in conjunction with the LOG or RDS options. If the ISC link fails or the RDS incurs a write error, the active IMS writes log records to indicate the failure. Failure of the alternate IMS to receive signals on the ISC link or the RDS is possible cause for failure; failure of the alternate IMS to receive new records on the log is not. For example, suppose you specify SURV=LNK,LOG, and the alternate IMS stops receiving signals on the ISC link. The failure of signals over the ISC link indicates a takeover, but as long as the log continues to send records satisfactorily, IMS does not request a takeover. In this case, IMS assumes that the ISC link itself is experiencing difficulty and is not a reliable indicator.

Do not confuse the reading of the log with LOG as a surveillance mechanism by the alternate IMS. XRF IMS requires that the alternate IMS read the log; XRF IMS depends on information in the log to update its control blocks. Only when you specify SURV=LOG does the alternate IMS periodically check the log to ensure that the active IMS is sending new records. Failure of the alternate IMS to receive new records does not cause a takeover. IMS uses the LOG option to confirm or veto a decision that the alternate IMS might make based on the lack of signals from the RDS or the ISC link. Do not use the LOG option as the only surveillance mechanism; it is best used to validate that the system is indeed not working if the LNK or RDS fails. Absence of LOG records might simply mean that no work needs to be done during the specified interval.

The active IMS writes to the log and the alternate IMS reads the log at two different rates. If the alternate IMS is not caught up, it reads the log continually. When it is caught up, the interval value on the LOG parameter controls the frequency with which the alternate IMS reads the log. Therefore, this interval can affect the level of I/O activity. One way to reduce the work of the alternate IMS is to increase the interval you specify on the LOG parameter.

If you set the interval too low, I/O activity to data sets increases, which can affect the performance of the active IMS.

The RDS is an important surveillance tool. If the ISC link fails or the system is inactive and no log records are written, the RDS continues to be updated. For example, if you specify SURV=LNK,RDS and z/OS enters a wait state, the ISC link and the RDS cease sending periodic signals. This warns the alternate IMS of a problem. If you specify no surveillance, the alternate IMS does not receive warning of the wait state. In that case, an operator must notice this condition and request a takeover. The following table shows how the alternate IMS uses the RDS, the IMS link, and log records to learn of certain events.

*Table 50. Informing the alternate IMS of events in the active IMS*

| Event                              | Log record tells the alternate IMS | Surveillance option that informs the alternate |
|------------------------------------|------------------------------------|------------------------------------------------|
| IMS abend                          | Yes                                | None                                           |
| z/OS abend                         | No                                 | LNK, RDS                                       |
| z/OS loop <sup>1</sup>             | No                                 | LNK, RDS                                       |
| z/OS wait state <sup>1</sup>       | No                                 | LNK, RDS                                       |
| CPC failure                        | No                                 | LNK, RDS                                       |
| VTAM failure that takes TPEND exit | Yes                                | None                                           |
| VTAM wait                          | No                                 | LNK,RDS                                        |
| IRLM failure                       | Yes                                | None                                           |

**Note:**

1. Surveillance-detectable loops and waits.

Although using the LNK and the LOG as surveillance requires no effort beyond specifying the initial control statements, you might not have an ISC link already in place. In this case, establish the ISC link through the 37x5 Communication Controller that controls your class-1 terminals.

If you do not specify the SURV parameter, IMS uses all three options: LNK (the ISC link), the RDS, and LOG.

The default frequency of signals from the active IMS for the ISC link is 3 seconds and for the RDS is 1 second. Two factors to consider in setting the timing for LNK and RDS are the speed of communication over the ISC link and the performance overhead on z/OS in the alternate IMS.

The default frequency for the alternate IMS to check for new records in the log is 1 second.

**RSENAME=**

The recoverable service element (RSE) consists of the active and alternate subsystems. Although failures of VTAM, z/OS, and the CPC can cause a takeover, these elements of the XRF complex depend on IMS to recognize a failure and initiate a takeover. The RSENAME is a 1- to 8-character ID. It is through this ID that IMS is known to DBRC, IRLM, and the z/OS availability manager. Signon and authorization of DBRC databases is done using the RSENAME keyword. RSENAME is the ID used to notify AVM of the active and alternate subsystems. It is also the name the operator sees in AVM messages. RSENAME is required for XRF.

**MNPS=**

The MNPS parameter defines the MNPS ACB for IMS systems in an XRF

complex that uses MNPS. This name must be the same in the procedures of both the active IMS system and the alternate IMS system. When users log on to the XRF complex, they enter this name as the IMS logon name. When a takeover occurs, the instance of the MNPS ACB on the active IMS is closed and the alternate IMS opens a new instance using the same MNPS ACB name. Specifications made using this keyword override any USERVAR specifications.

#### **MNPSPW=**

Use this keyword to define a password for the MNPS ACB of the IMS systems in an XRF complex. Each IMS system provides this password to VTAM when it opens an MNPS ACB. You must also set VTAM to check for passwords. If VTAM expects a password and you have not specified one, the MNPS ACB is not opened.

#### **USERVAR=**

The USERVAR parameter defines the USERVAR for IMS in an XRF complex that uses USERVAR. The USERVAR is stored in both a USERVAR table and an interpret table. The USERVAR is linked to the APPLID of the active IMS system in the USERVAR table. The USERVAR is linked to a logon message in the interpret table. Users logging on to the XRF complex need only know the logon message to be connected to the active IMS system. When a takeover occurs, IMS places the APPLID of the new active IMS system in the USERVAR table of the active IMS system.

If a value is entered in the MNPS parameter, the USERVAR parameter is overridden.

#### **SWITCH=**

The following list describes the options on the SWITCH parameter and describes how the alternate IMS uses the option in considering a takeover.

##### **Option Alternate IMS considers a takeover when:**

**IRLM** A log record tells it that IRLM has failed.

**LNK** Signals on the ISC link fail to appear after a specified time.

**LOG** New records on the log fail to appear after a specified time.

**RDS** Time stamps on the RDS link fail to appear after a specified time.

##### **TPEND**

A log record tells it that VTAM has taken an IMS TPEND exit and failed.

To establish failure of VTAM as a cause for takeover, specify the TPEND option on the SWITCH parameter.

If you use IRLMs to manage local locks or to manage the resources the XRF complex shares with other IMS systems, you might want failure of IRLM to be a cause for takeover. Specify the IRLM option on the SWITCH parameter. If you want takeover on IRLM, specify AUTO=YES.

Specify the surveillance mechanisms on the SWITCH parameter that are critical in alerting the alternate IMS of certain failures. You can combine these options and use them more than once. For example, if you specify SWITCH=(LNK,LOG), IMS considers a takeover if the ISC signal fails to appear and new log records fail to appear after a specified time. One of these events without the other does not cause IMS to consider a takeover.

Specify a timeout number on the LNK, LOG, and RDS parameter that tells the alternate IMS how long it should wait for a signal before considering a

takeover. The default timeout is 9 seconds for the ISC link and 3 seconds for the LOG and the RDS. These times are only approximate; actual values might be slightly larger.

The default for the SWITCH parameter is:

SWITCH=(LNK,LOG,RDS),(TPEND),(IRLM)

The SWITCH parameter tells the alternate IMS to consider a takeover if VTAM fails and if the result is an IMS TPEND exit, or if IRLM fails and the result is an IMS STATUS exit, or if all three surveillance mechanisms alert the alternate IMS of problems.

When you specify the timeout interval, consider the following exceptional conditions:

- Software or hardware problems might cause the system to stop processing IMS work for 30 to 60 seconds.

Certain software or hardware problems might cause wait states or loops that last up to a minute in duration. Examples of these kinds of problems are software recovery routines, hardware spin loops, or hardware recovery procedures. If your timeout interval is set at the default of 9 seconds for LNK, 3 seconds for RDS or 1 second for LOG, these conditions cause a takeover. If the problem becomes chronic, you must identify the cause and make the necessary changes. Before you make the correction, you can either:

- Raise the timeout intervals on the LNK, LOG, and RDS parameter statements so that the condition does not cause a takeover
  - Specify NO on the AUTO parameter to allow the operator to take control at takeover, giving the condition more time to correct itself and possibly avoid a takeover
- A SLIP trap specifies that a restartable wait state be entered when the trap conditions are met. This stops the active IMS long enough for the alternate IMS to take over. The re-startable wait state causes a takeover. It seems likely that a takeover is appropriate in this situation. When the trap occurs, a lengthy analysis period and possibly a stand-alone dump follows.
  - An operator stops a CPC

Surveillance should be stopped when the operator wants to stop the CPC. The operator might do this when entering new JES, because of an SDUMP or a disabled console. If your operator stops the active CPC and has not stopped surveillance, the lack of signals from surveillance mechanisms causes the alternate IMS to request a takeover. If you do not want a takeover to occur, discontinue surveillance or specify AUTO=NO to give control to an operator before a takeover proceeds.

Because takeover can happen when a set of time-dependent parameters are satisfied, setting up surveillance intervals is an important administrative task. Also, those intervals might not be suitable for all workloads run by the active IMS system.

Factors that affect surveillance intervals are:

- CPC power or type (for example, UP or MP)
- Operational procedures and response objectives
- Recovery/delay situations where a timeout is exceeded
- For critical work, the relative tolerance to unplanned outage

#### **ALARM=NO | YES**

This parameter requests that a service processor notify the operators when a takeover begins. The default is NO.

When an automatic takeover of the active IMS by the alternate IMS occurs, the service processor alarm on the alternate IMS sounds if YES has been specified. The alarm also sounds when operator intervention is needed to manually switch the system.

#### **AUTO=YES | NO**

AUTO establishes whether the takeover proceeds automatically when IMS requests it or whether the operator intervenes. If you specify AUTO=YES, a takeover occurs automatically (based upon a failure in the host as detected by the surveillance mechanism). If you specify AUTO=NO, the operator must manually force the system to switch. Message DFS3869 appears on the master console of the alternate IMS when the alternate IMS detects, through the surveillance mechanism, that a takeover is indicated. When AUTO=NO is specified, the operator can delay the takeover while trying to resolve the problem on the active IMS. Normal tracking continues on the alternate IMS while the problem on the active IMS is being fixed. The takeover message ceases to be sent to the active IMS after the problem is corrected or after the operator initiates the switch.

Operator intervention undoubtedly extends the time before NCP switches sessions on XRF terminals and IMS reestablishes sessions on terminals. Operator intervention also causes disruption of the workload processing. Your installation, though, might want the operator to perform certain tasks before the takeover proceeds. The tasks might be:

- Ensuring that the takeover decision is valid
- Manually switching terminals or DASD
- Quiescing of jobs on the alternate IMS

If it is important to control the quiescing of the work in the alternate IMS, request that the takeover not proceed without operator intervention. This gives your operator time between the takeover request and the takeover to cancel the jobs.

If you do not want the operator to control the takeover, specify AUTO=YES. The default is AUTO=NO.

#### **KEYEVENT=MSG | NONE**

KEYEVENT determines whether your IMS operators receive all of the IMS messages for XRF. Although the takeover messages must be displayed, you can choose to have certain messages suppressed. When you specify KEYEVENT=MSG, all messages are displayed. When you specify KEYEVENT=NONE, optional messages DFS3882 through DFS3888 are suppressed. The default is NONE.

#### **DEFERFIX=xx**

This parameter indicates the DFSFIXxx PROCLIB member that should be processed if page fixing of non-Fast Path resources on the alternate IMS is to be deferred until takeover. IMS ignores Fast Path page-fixing options at this time. For the alternate IMS, only those DL/I blocks and routines that can be page fixed with the existing DFSFIXxx IMS.PROCLIB member are page fixed.

#### **DFSFIXxx**

The DFSFIXxx PROCLIB member specifies the page-fixing values for an IMS system. In an XRF complex, two members can be used. One member specifies what is to be page fixed at initialization time for both the active and alternate IMS systems. This member is pointed to by the FIX=xx parameter of the startup procedure. The second member is pointed to by the DEFERFIX=xx parameter in



DFSHSBxx member of IMS.PROCLIB. For the active IMS, page fixing requested in this member occurs at initialization. For the alternate IMS, page fixing requested in this member occurs at takeover.

Page fixing done on the alternate IMS during initialization speeds the takeover but takes additional real storage on the alternate IMS throughout the tracking phase.

## DFSVSMxx

The DFSVSMxx PROCLIB member contains the log data set definition information, specifies the allocation of OLDS and WADS, the number of buffers to be used for the OLDS, and the mode of operation of the OLDS (single or dual). Dual OLDSs are recommended, because when loss of the OLDS occurs and no backup is available, the XRF complex fails. Because the alternate IMS allocates the log data sets defined to it and any data sets used by the active IMS, the mode of operation and buffer definitions must be the same for both subsystems. Both active and alternate IMS systems must use the same member.

## Defining the VTAM USERVAR table

These examples show the relationship between the IMS system definition parameters and the VTAM interpret USERVAR table parameters used with Extended Recovery Facility (XRF).

### Related reading:

- For more information on the USERVAR table, see VTAM operations in an XRF complex with USERVAR.
- For more information on IMS system definition parameters, see *IMS Version 11 System Definition*.

You should perform the following tasks:

- You must specify the APPLIDs and passwords of both systems in the APPLID and PASSWORD parameters of the COMM macro at IMS Generation. For example:  
`IMSCTRL HSB=YES COMM APPLID=(IMS1,IMS2),PASSWORD=(IMSP1,IMSP2)`
- Before starting the IMS system, you must set up the DFSHSBxx member of the IMS-PROCLIB specifying the Generic Application ID in the USERVAR parameter. For example:  
`USERVAR=IMS`
- The Generic Application ID must be registered into the VTAM interpret table as a user variable. For example:  
`LOGCHAR APPLID=(USERVAR,IMS),SEQUENCE='IMS'`
- APPLIDs specified in APPLID parameter of the COMM macro must be defined as either SPO (secondary programmed operator) or PPO (primary programmed operator) in the VTAM definition library. For example:  
`VBUILD TYPE=APPL IMS1 APPL AUTH=(ACQ,SPO),PRCT=IMSP1,HAVAIL=YES  
IMS2 APPL AUTH=(ACQ,SPO),PRCT=IMSP2,HAVAIL=YES`

## Placement of IMS data sets in the XRF configuration

Although IMS running with an alternate IMS in another CPC is managed operationally as a single system, you do need to plan in detail for duplication of IMS system data sets.

The three main considerations for placing your data sets are:

- Availability of data sets during tracking and takeover

An XRF complex consists of two systems that must sometimes access the same data sets or identical copies of the same data sets. Therefore, IMS requires that you place some data sets on DASD shared by the two systems.

**Recommendation:** Place other data sets on shared DASD. However, you can switch some data sets through a switching device or maintain separate copies of them.

- Prevention of single points of failure

IMS requires that you maintain (and constantly synchronize) separate copies of some data sets for the two systems.

**Recommendation:** Maintain separate copies of other data sets.

- Accessibility of data sets to one IMS system

**Recommendation:** Keep the data sets that are unique to one system on local (that is, non-shared) DASD.

For the best performance at takeover, keep the shared and non-shared I/Os separate from each other. For example, place the IMS database data sets on different volumes, controllers, and channels than the z/OS data sets.

Placement of the system logs is important. A critical step in the takeover process is the isolation of the databases from the failing active IMS. When the active IMS and the alternate IMS run on different CPCs, the alternate IMS can stop the failing active IMS from accessing the OLDS and WADS by reserving the DASD devices on which they reside. Therefore, do not place other data sets on the DASD that contain the OLDS and the WADS.

You should dynamically allocate all IMS databases and area data sets. If they do not already exist, you must generate the DFSMDA members for all IMS databases and areas.

**Related reading** For more information on generating the DFSMDA members for all IMS databases and areas, see *IMS Version 11 System Definition*. All IMS Full Function database names and DEDB area names must be unique.

Some of the requirements and recommendations in this topic address the protection of your data sets. Always protect your data sets to the extent your resources allow. The following lists summarize the required and recommended placement of these data sets.

The following list shows the recommended placement of data sets and data in database A (associated with the active IMS):

- IMS.LGMSG
- IMS.LGMSGSL
- IMS.SHMSG
- IMS.SHMSGSL
- IMS.QBLKS
- IMS.QBLKSL
- IMS.MSDBDUMP
- SYSDUMP data
- SYSABEND data



The following list shows the recommended placement of data sets and data in database B (associated with the alternate IMS):

- IMS.LGMSG
- IMS.LGMSGSL
- IMS.MSDBDUMP
- IMS.QBLKS
- IMS.QBLKSL
- IMS.SHMSG
- IMS.SHMSGSL
- SYSDUMP data
- SYSABEND data

The following list shows the recommended placement of data sets and data in databases C and D (associated with both the active and alternate IMS systems):

- IMS.ACBLIBA
- IMS.ACBLIBB
- IMS.FORMATA
- IMS.FORMATB
- IMS.JOBS
- IMS.MODBLKSA
- IMS.MODBLKSB
- IMS.PGMLIB
- IMS.PROCLIB
- IMS.SDFSRESL
- IMS.TFORMATA
- IMS.TFORMATB

### **Additional data sets required for XRF**

The following figure illustrates the information from the above lists, and shows you the connections between the four databases (A-D), the active and alternate IMS systems, and the shared data sets.




Figure 79. Recommended data set placement

The additional data sets listed in the figure are present in control region JCL in an XRF complex.

Table 51. Additional data sets required by XRF

| Data set  | Use                       |
|-----------|---------------------------|
| IMS.RDS2  | Separate restart data set |
| IMS.LGMSG | Local message queue       |

Table 51. Additional data sets required by XRF (continued)

| Data set     | Use                         |
|--------------|-----------------------------|
| IMS.SHMSGL   | Local message queue         |
| IMS.QBLKSL   | Local message queue         |
| IMS.MODSTAT2 | Separate MODSTAT data set   |
| IMS.MSDBCP3  | Alternative MSDB checkpoint |
| IMS.MSDBCP4  | Alternative MSDB checkpoint |

## Data sets that must be shared in the XRF complex

In addition to critical data sets that make up databases, several system data sets need to have a single copy and be shared by both the active and alternate subsystems. These have the following DD names:

- RECON1, RECON2, RECON3
- DFSOLPnn
- DFSOLSnn
- DFSWADSnn
- IMS Spool data sets
- IMS.MODSTAT
- IMS.MODSTAT2
- IMS.RDS
- IMS.RDS2
- IMS.MSDBCP1
- IMS.MSDBCP2
- IMS.MSDBCP3
- IMS.MSDBCP4

**Recommendation:** For added efficiency and ease of operation, the DL/I databases should be on DASD shared by the two CPCs. You might, however, choose to switch these data sets through a switching unit, such as an IBM 3814, or a channel-switch on an IBM 3880 Storage Control.

It is also recommended that you have two copies of the OLDS shared by the two CPCs. If you have one copy of the OLDS and a permanent I/O error or power failure causes physical loss of the one copy, tracking cannot continue. In this case, a failure of the new active IMS occurs.

Figure 79 on page 659 identifies the data sets that you should place on shared DASD. In the figure, they are on the DASD located between the active IMS and the alternate IMS. Notice that the system logs, the restart data set (RDS), MODSTAT, and MSDB data sets are not only on shared DASD, but separate copies also reside on shared DASD. The active IMS maintains these copies.

## Data sets not shared in the XRF complex

For best performance, place data sets that contain information unique to one system on local, non-shared DASD, and define these data sets in separate catalogs. The data sets you maintain separately contain information unique to one system, such as the data a system needs when it attempts to recover from a failure.

Figure 79 on page 659 shows these data sets on DASD labeled A and B. They could also be on DASD that is shared by the two systems.

Maintain certain data sets as distinct library copies that are not to be shared. This is to reduce the points of failure. These system data sets are:

- IMS.ACBLIBA(B)
- IMS.FORMATA(B)
- IMS.TFORMATA(B)
- IMS.MODBLKSA(B)
- IMS.JOBS
- IMS.PROCLIB
- IMS.SDFSRESL

### **Data sets that must be duplicated**

Some data sets must also be separate data sets. Those system data sets are:

- IMS.MSDBDUMP
- IMS.LGMSG
- IMS.SHMSG
- IMS.QBLKS
- IMS.LGMSGSL
- IMS.SHMSGSL
- IMS.QBLKSL
- IMS Spool data sets
- SYSUDUMP and SYSABEND data sets

### **Making online changes in an XRF complex**

You can still use the Online Change utility in the XRF complex. When you use it to maintain separate but duplicate copies of IMS data sets, make the identical changes to both copies of a data set. Then issue the /MODIFY PREPARE, /DISPLAY MODIFY. and /MODIFY COMMIT commands at the active IMS.



---

## Chapter 45. Remote Site Recovery overview

This topic describes the Remote Site Recovery (RSR) complex, explains how to install RSR, compares RSR to XRF, explains how to initialize RSR, and describes IMS error handling for RSR.

### Related concepts:

“The RSR environment” on page 10

 Recovery of IMS using Remote Site Recovery (RSR) (Operations and Automation)

---

### RSR overview

RSR allows you to recover quickly from an interruption of computer services at an active (primary) site. RSR supports recovery of IMS DB full-function databases, IMS DB Fast Path DEDBs, IMS TM message queues, and the IMS TM telecommunications network.

When your computing system is disabled, you need to recover quickly and ensure that your database information is accurate. Interruption of computer service can be either planned or unplanned. When interruption on the primary computing system occurs, you need to resume online operations with minimal delay and minimal data loss.

IMS database and online transaction information is continuously transmitted to a tracking (remote, or secondary) site. The remote site is continually ready to take over the work of the active site in the event of service interruption at the active site.

Because customers require that IMS has the ability to resume online operations at a remote site in the event of an extended outage (either planned or unplanned) at the active site, RSR does the following:

- Provides a remote copy of the necessary IMS DB and IMS TM log records for database and message queue recovery at the remote site.
- Reduces the time required to resume computer service to usually less than an hour.
- Allows you select and filter out the log records that are not needed to support the defined critical environment.
- Continues to operate when the active or remote sites or the RSR transmission facility become temporarily unavailable, and provides a way to resynchronize the sites as soon as possible.
- Provides transaction consistency between the active and remote sites.
- Supports IMS DB and IMS DBCTL. Supports full-function databases and Fast Path DEDBs. Supports both online IMS DB and DBCTL workloads, as well as batch workloads, at the active site.
- Supports coordinate disaster recovery processing for IMS and DB2 for z/OS.
- Supports data sharing at the active site.
- Coexists with XRF (see Table 55 on page 680).
- Recognizes that DBRC is operating at the active site and, separately, at the remote site.

- Supports standard ACF/VTAM communication protocols, so that new technology is not required for data transmission.

Remote Site Recovery at the active site provides the following general functions:

- Allocation, control and termination of conversations between active IMS loggers or the isolated log sender and the tracking IMS log router
- Sending of log data to the remote site; previous missing log data is sent independently of current log data
- Support for the transition of an active site to assume the role of the remote site after a remote takeover

Remote Site Recovery at the remote site provides the following general functions:

- Allocation, control, and termination of conversations between tracking IMS log router and the active IMS loggers and the isolated log sender; any errors in log data transport are recognized and corrected when possible
- Receipt of log data and writing it to tracked SLDs
- Collection of database change log data
- Application of changed data to shadow databases, if necessary
- Maintenance of the MODSTAT data sets to indicate current ACB, FORMAT, and MODBLKS libraries in use
- Maintenance of the remote site's RECON data set, including status of databases and tracking and shadow logs, as well as active site information and remote site information
- Monitoring of an XRC session to support coordinated disaster recovery for DB2 for z/OS and IMS
- Support for tracking IMS operations, including starting, tracking, terminating, restarting, and abnormally terminating

**Note:** The tracking IMS does not support transaction processing or the running of dependent IMS regions while in tracking mode.

- Support for a subset of IMS utilities at the remote site while in tracking mode
- Support for the transition from remote site to “new active” site after remote takeover

The following table shows which types of IMS systems can be used together in an RSR environment. A DB/DC IMS can track all three types of IMS systems: DB/DC, DBCTL, and DCCTL. A DBCTL IMS can track DB/DC or DBCTL subsystems. A DCCTL IMS can only track DCCTL subsystems.

*Table 52. Combinations of IMS types in an RSR environment*

| Active IMS type | Type DB/DC tracking IMS | Type DBCTL tracking IMS | Type DCCTL tracking IMS |
|-----------------|-------------------------|-------------------------|-------------------------|
| DB/DC           | Yes                     | Yes                     | No                      |
| DBCTL           | Yes                     | Yes                     | No                      |
| DCCTL           | Yes                     | No                      | Yes                     |
| DB Batch        | Yes                     | Yes                     | No                      |
| TM Batch        | Yes                     | No                      | Yes                     |

#### **Related concepts:**

“Recovery with a Remote Site Recovery system” on page 114

---

## Requirements for using RSR

The first and most important requirement for using RSR is that you must have a remote site available. In order for the remote site to be able to take over your critical workload from the active site, the remote site must have almost everything the active site has, including hardware, software, and at least some personnel.

Before you use RSR, you need to consider RSR's effect on the following tasks:

- Installation
  - Installing IMS with RSR at two, usually physically separate, sites.
- Operations
  - Maintaining a remote site.
  - Establishing new procedures for:
    - Database recovery
    - System recovery
    - Network switching
- Database administration
  - Considering the DASD space requirements for the remote site.
  - Considering the database recovery activities, such as:
    - Deciding the frequency of recovery: Do you recover only in case of disruption, or continually?
    - Synchronizing remote and primary copies (such as maintaining consistency across reorganization).
    - Sending image copies to the remote site.
  - Maintaining the remote supporting environment, such as PSBs and DBDs.
- System programming, maintenance, and tuning
  - Writing exit routines for new functions.
  - Implementing new log structure and format.
  - Maintaining the remote supporting environment (including system definition and libraries).
  - Preparing for system and database recovery (IMS restart).
  - Writing new audit procedures.
- Network administration
  - Examining your current network (VTAM) configuration and assess how the workload added by RSR will affect it.

---

## Basic components of RSR

To provide Remote Site Recovery for your IMS, you need a primary IMS site and a remote IMS site. The primary site is where your IMS work is performed. The remote site performs no IMS work, but provides remote recovery support and is ready to perform your IMS work at any time.

### IMS systems

Active IMS systems and XRF alternate IMS systems are a part of the *active site*. A tracking IMS is at the *remote site*.



An *active IMS* is the subsystem used to perform your day-to-day work. It can consist of a DB/DC, DBCTL, or DCCTL system, as well as batch environments. The active IMS supports both databases and terminals. Databases at the active site are called *master databases*.

An *alternate IMS* is an online IMS located at the same site as the active IMS. It is the XRF alternate IMS for the active IMS, and is therefore used only if you are using XRF at the active site. The alternate IMS uses the master databases of the active IMS.

The *tracking IMS* is usually located at a different site from the active and alternate IMS systems. This is the IMS that tracks the active site activity by maintaining backup copies of the log data of the active IMS. It can also make copies of the master databases from the active IMS. These backup databases are called *shadow databases*.

The tracking IMS can only be an online IMS. Batch environments are not supported as tracking IMS systems.

## The Transport Manager Subsystem

The transport manager subsystem (TMS) provides communication services to components of RSR, using VTAM's advanced program-to-program communication (APPC) support. The TMS:

- Dynamically allocates APPLIDs to IMS systems, as needed.
- Provides directory support, allowing RSR components to run on any CPC in your installation without system definition changes. This directory support also allows RSR IMS systems to dynamically locate IMS systems operating in particular roles (for example, as active or tracking).
- Provides full-duplex conversations. These conversations do not have a send or receive state, so each end of the conversation can send or receive at any time and can do so simultaneously.
- Provides a single service address space for the isolated log sender per CPC.
- Provides a simple set of interfaces for RSR components.

The RSR TMS provides an interface similar to APPC, and includes capabilities that RSR requires when APPC is inappropriate. Because a TMS conversation is full duplex (using two APPC conversations), it eliminates much of the state management complexity of APPC and provides better performance in areas such as bandwidth and CPC usage.

## The log router

The log router component of the tracking IMS receives data from the active IMS systems, stores the log data in tracking log data sets, and routes log records to individual tracking subcomponents, called *tracking IMS systems*. The log router is unique to tracking IMS systems; it is not found in active IMS systems.

**Note:** In the rest of this topic, the term “SLDS” (system log data set) refers to the active-IMS-generated SLDS stored and managed by the tracking IMS at the remote site.

The log router manages the tracking end of communication between the active and remote sites. It initiates and accepts conversations with the active IMS logger.

These conversations enable the log router to receive log data and to communicate system control information from the active IMS logger to the log router component.

The log router receives log data from the active IMS loggers and writes the log data to SLDS data sets. When the log router creates an SLDS, it uses data set allocation parameters found in the IMS PROCLIB member DFSRSRxx. If more than one active logger sends log data to the tracking IMS, data reception and log data writing proceeds in parallel. Later, you can copy the SLDS to other data sets (such as tape), and delete the original SLDS; or you can have the log router perform automatic SLDS archiving.

When the tracking IMS operates at the database readiness level, log records are presented to the database tracking IMS systems (DL/I and Fast Path) after they are written to an SLDS. If any active IMS systems participate in block-level data sharing, the log records from the sharing IMS systems are merged in time stamp sequence before being presented to the database tracking IMS systems.

When an active IMS is unable to send log data to the tracking IMS (because of a repairable cause, such as network communication disruption) and the log router recognizes a gap in received log data, the log router contacts the isolated log sender at the active site and requests the missing log data. After the data is received, it is written to the SLDS and presented to the tracking IMS systems. Until the log gap is filled, no log records beyond the gap are sent to the tracking IMS systems. The records after the gap are written to the SLDS. Later, these records will be presented to the tracking IMS systems in order to bring them up to date with the latest possible active IMS log data. This procedure is called *catch-up processing*.

A stopped shadow database or area can be restarted at the remote site, at which time the database or area will be made to match the active site with online forward recovery. *Online forward recovery* (OFR) is the process of bringing data in a database or area to a current state. The log router reads log records from previously written SLDSs and presents them to the tracking IMS systems. Eventually, the log read process for OFR catches up with the current log routing and all further database updates are handled by normal tracking.

The log router records (on the tracking IMS log) its location in each SLDS and records the log records that have most recently been routed to the tracking IMS systems. This information is used when the tracking IMS restarts and enables log records to be routed to the tracking IMS systems in the correct order.

**Related reading:** For more information on database readiness levels, see “Determining the extent of recovery” on page 672.

## Isolated log sender

Certain conditions exist that prevent IMS log data from being sent to the remote site at the time it is written to the active IMS's logs. These conditions might be VTAM link failures, a temporary outage of the tracking IMS, or other repairable problems.

When the active IMS is unable to send its log data to the tracking IMS, but is still logging database changes, the tracking IMS recognizes that a *gap* exists in the log. When the tracking IMS is able, it requests the missing log data—the gap—to be sent from the active IMS. It is the isolated log sender (ILS) that sends the missing log data to the tracking IMS.

The ILS is functionally active only at the active site, but should also be available at the remote site so that the ILS can be activated after remote takeover. Only one ILS is functionally active at a time for any global service group sending missing log data to the tracking IMS, but standby ILS instances can also be started. The ILS can be started during TMS initialization or by an operator command.

Because it is possible for a log to have a gap that spans several log data sets, the ILS can send the log data sets in parallel in concurrent conversations with the tracking IMS.

Until the ILS fills the gap for the tracking IMS, all subsequent log data is simply written to the log data set of the tracking IMS and is not used by the database tracking IMS systems. This unused log data is *marooned* until the gap is filled. The following figure illustrates the log gap.




Figure 80. Marooned log data

## DL/I database tracking IMS

The DL/I database tracking IMS receives the log data from the active IMS from the log router and the tracking IMS updates shadow databases at the remote site. This tracking IMS is active only when you choose database readiness level (see “Determining the extent of recovery” on page 672) and LSO=S is specified in the EXEC parameters.

The tracking IMS requires exclusive authorization for using the database before updates can be applied. It also checks with DBRC to determine if the updates should be applied. The IRLM is not required at the remote site for database tracking.

The log router calls the database tracking IMS periodically, to ensure a short restart time for the tracking IMS and to record a pointer into the logs being processed to allow correct reprocessing from a particular point in the logs. These points are called *milestones*. Milestones occur periodically through the life of a tracking IMS, and are triggered at regular intervals. A new milestone does not begin until the previous one completes. When the log router calls the database tracking IMS, shadow database updates are written to DASD.

Database tracking also includes an online forward recovery function. This allows catch-up database processing so that databases that were unavailable for some reason can be made current. The database tracking IMS also uses online forward recovery to update databases when an initial or current image copy is received.

Online forward recovery is only available on tracking IMS systems running at database readiness level (DLT) for databases defined as DBTRACK (database level shadowing).

## Fast Path database tracking IMS

The Fast Path database tracking IMS receives the Fast Path log data from the log router of the active IMS. This tracking IMS is active only when you choose database readiness level (see “Determining the extent of recovery” on page 672). The IRLM is not required at the remote site for database tracking.

The FP database tracking IMS supports all the availability features of DEDBs for shadow databases: database partitioning (areas), multiple copies (MADs), and records deactivation (EQEs). The number of area data sets at the tracking IMS can be different than that at the active site. You must register each area you want tracked in RECON as database-level shadowing.

MSDBs are not supported by the Fast Path database tracking IMS. VSO DEDBs are handled in the same way as non-VSO DEDBs.

As Fast Path database updates are received by the Fast Path database tracking IMS, they are stored in a z/OS data space to reduce physical I/O. The database updates are kept in the data space until either a data space threshold value is reached or the tracking IMS writes the data to disk. If several updates apply to the same CI, they are accumulated and written to DASD at the same time.

DEDB update log records are held until a commit log record appears. This is done because Fast Path does not log UNDO log records (which DL/I uses to be able to back out changes).

## Naming conventions

In order to uniquely identify the various parts of an RSR complex, an RSR name is qualified according to a naming convention. As part of this convention, various parts of the RSR complex are subdivided and named, so that each unique part of the RSR complex is given a unique name.

A fully qualified RSR name would look like this:

GSGname.SGname.SYSTEMname.INSTANCEname.COMPONENTname

The parts that make up this name are described in the following list.

### Global service group

A *global service group* (GSG) is the collection of all IMS systems that access a particular set of databases. A global service group can span several z/OS subsystems at more than one geographical location.

You can have more than one global service group. IMS systems that do not share resources are generally in separate GSGs, but are not required to be. Systems belonging to different GSGs can be connected using MSC or ISC links and can cooperate by transaction routing.

### Service group

A *service group* (SG) is a collection of all IMS systems that access a particular version of a GSG's resources, including the recovery control

(RECON) data set. A service group usually includes one or more subsystems at a single site, with the databases and the RECON data set shared between the subsystems.

An RSR GSG is made up of two service groups: the active service group and the tracking service group. A service group name is usually the site name. When a remote takeover is required, it occurs on a service-group level: All activity of one service group is taken over by another service group.

The following figure shows two service groups in a very simple RSR complex. The two service groups—the active site and the remote site—make up the global service group. The active site contains the active IMS, its database, log, and RECON data set. The remote site contains the tracking IMS, its database, log, RECON data set, and the active IMS's SLDS.




Figure 81. RSR service groups and global service groups

### System

For RSR, a “system” is a z/OS implementation that has one instance of the RSR TMS running on it. Because of the need for operational flexibility, the system name can be generated dynamically.

### Instance

For RSR, an “instance” is a particular, unique IMS subsystem within the z/OS system. An example of an instance could be a particular DBCTL subsystem, a particular DB/DC subsystem, a batch DL/I job, or a batch utility job. For DB/DC, DBCTL, and DCCTL subsystems, the IMSID is the

instance name. Several IMS subsystems running on one z/OS system have the same system name but different instance names.

### **Component**

For RSR, a “component” is a part of RSR within a particular IMS instance. Examples of components are the logger, log router, and isolated log sender.

## **Remote takeover**

Because “takeover” is such an integral part of RSR, this topic defines what is meant by the term in the RSR context. Using RSR, you have the capability to transfer an installation's IMS processing capacity to a remote location from an active site if the need arises. This transfer of processing responsibility is called a remote *takeover*. RSR provides two types of takeovers: planned and unplanned.

A planned takeover, initiated by the operator, is an orderly transfer of operation to the remote site after a successful shutdown at the active site.

An unplanned takeover is the switching to a remote site as a result of an unexpected outage at the active site. This takeover, also operator-initiated, does not require a successful active shutdown. For both takeover types, however, a successful shutdown of the tracking IMS is required. This shutdown is initiated by a takeover command. The operator then restarts the active IMS systems at the remote location.

---

## **RSR processing**

IMS recovery processing, as well as user tasks and procedures, are unchanged, except that recovery is extended to the remote site.

- The IMS log continues to be one of the basic elements in database and subsystem recovery. System log data is archived when it is sent from the active site to the remote site. This data can be processed as it is received to keep copies of databases updated.
- Database Recovery Control (DBRC), required for subsystems using remote recovery, continues to perform central recovery control functions.  
DBRC provides the following functions for RSR:
  - Commands for defining, updating, and displaying RSR status
  - Services for the active IMS to determine what remote site is used, and which databases are associated with that site
  - Facilities for the remote site to record information about log data sent from the active site and received at the remote site
  - Services to assist in takeover
- You can use emergency restart to ensure consistency of databases and to recover resources such as message queues and scratchpads.

---

## **Determining the extent of recovery**

After you define the active and remote sites, you must decide what level of readiness you need. RSR provides two levels of readiness for recovery: recovery readiness and database readiness.

You control which readiness level you want in two ways:

- Use the IMS RLT (recovery level tracking) or DLT (database level tracking) startup options to determine the readiness level of the tracking IMS.

RLT is a valid readiness level for any type of tracking IMS system: DBCTL, DB/DC, and DCCTL.

#### **Recovery level tracking (RLT)**

DLT is not allowed for a DCCTL tracking IMS system. A tracking IMS system can have a DLT readiness level only if it is DBCTL or DB/DC.

For RLT, you first send copies of the master databases from the active site to the remote site. Ensure that your databases are registered with DBRC. These copies of the master databases are not updated by the tracking IMS after they arrive at the remote site; that is, they are not shadow databases. The log data sent from the active site to the remote site is archived on tracked SLDSs and kept until recovery or takeover is required.

Because the databases must be forward recovered before you use them (after a takeover), you should use recovery level tracking only for those subsystems for which recovery is infrequent or which do not require quick recovery. In most cases, a remote takeover at recovery readiness level takes several hours.

Because the databases are not shadowed, you do not require the DASD resources for the remote databases, and the tracking IMS requires fewer CPC resources to maintain the recovery readiness level for RSR. This resource savings allows the remote site to perform other, non-RSR, tasks while maintaining recovery readiness and tracking the active IMS.

**Recommendation:** Also, because the databases are not shadowed, before a planned remote takeover, you should change the readiness level from RLT to DLT, restart the tracking IMS, and start the databases (or areas) so that they can be automatically recovered in parallel using online forward recovery (OFR). You can periodically recover databases offline to shorten the recovery time after a remote takeover.

For an unplanned remote takeover, you must use the database recovery utility to recover the covered databases before starting the new active IMS.

#### **Database level tracking (DLT)**

For database readiness level (DLT), you first send copies of the master databases from the active site to the remote site. After the database copies are at the remote site, the user must NOTIFY.IC to register the IC in the remote site RECON data set. Then, the user can either do a GENJCL.RECEIVE to install the shadow database, or install the shadow database using some other user installation method and issuing a NOTIFY.RECOV command to indicate that the database is installed. When changes are made to the databases at the active site, those same changes are applied to the shadow databases. In this way, the shadow databases are always kept current, but because the remote site's process is asynchronous with the active site's, when a takeover occurs the shadow databases might not be absolutely current.

Because the databases from the active site are shadowed by the tracking IMS, takeover can usually occur in less than an hour.

Also because the databases are shadowed, you need to allocate permanent DASD resources for the databases being tracked.

You can choose to track only your most critical databases at the database (DBTRACK) readiness level and track your less critical databases at the



recovery (RCVTRACK) readiness level. This allows you to use fewer DASD and CPC resources at the remote site. You can change a database's tracking readiness level without having to restart either the active or the tracking IMS, but you must issue a database recovery (/DBR) command for the database, and you must change the database's readiness level at both sites.

You can change the readiness level of the tracking IMS (from DLT to RLT or from RLT to DLT) at any time, but when you do, you must restart the tracking IMS. The active IMS can be operational while you change the readiness level of the tracking IMS.

- Use the RCVTRACK (recovery readiness tracking) or DBTRACK (database readiness tracking) registrations of databases to track individual databases. RCVTRACK and DBTRACK define the readiness level for databases. If you define a database as either RCVTRACK or DBTRACK, that database is considered *covered*.

---

## Defining an RSR environment

A single active online IMS environment does not require a separate system definition for the tracking IMS. However, additions to the combined system definition for an RSR environment are necessary.

You need a separate tracking IMS system definition in the following cases:

- The tracking IMS is to track more than one active IMS, and no one system definition sufficiently defines all of the databases for the tracking IMS. You can eliminate this separate system definition requirement if you enlarge one of the active IMS systems to include the additional databases.
- Batch DL/I jobs use databases that are not defined to the active IMS.
- You want to eliminate definitions for non-tracked databases from the system definition of the tracking IMS.
- You want to reduce the data communications definitions from the system definition of the tracking IMS.

RSR allows you to provide separate master terminal operators for active, alternate, and tracking IMS systems at both the active and remote sites. You define the additional tracking IMS MTO and secondary master terminal by specifying MTOID=3 in the IMS procedure, and by specifying the three names in the COMM and TERMINAL macros. You can also use the APPLID $n$  override for the IMS /START command.

Because the roles of active and remote sites are reversed after a remote takeover, you should define both the active and tracking IMS systems to allow each to be easily restarted as either an active or tracking IMS.

The following parameters are very important to the proper configuration of an RSR TM environment:

- APPLID=(*name1,name2,name3*) (in the COMM macro)  
*name3* is used as the APPLID of the RSR tracking IMS. *name3* can match either *name1* or *name2*, but no two APPLIDs can be active in the network simultaneously.

You can override the defined VTAM application identifier for the IMS by using the APPLID $n$  parameter in the IMS procedure.

- APPLID1=*name1*,APPLID2=*name2*,APPLID3=*name3* (in the IMS procedure)

These can be specified in the startup parameters to override the APPLID names in the COMM macro. The requirement that the APPLIDs be unique makes their specification in the startup parameters highly recommended, especially if you want to run with common system definitions or common procedures.

Because APPLID1, APPLID2, and APPLID3 are specified in the COMM macro, they do not apply to a DBCTL tracking IMS.

- MNPS=*uname* (in the IMS procedure)

This optional parameter can be specified in the startup parameters. If specified for an RSR-capable active IMS system, the specification of this parameter must be the same at the remote site (after remote takeover).

For an XRF active IMS, this parameter overrides the XRF MNPS specified in the IMS procedure library member DFSHSBxx.

- USERVAR=*uname* (in the IMS procedure)

This optional parameter can be specified in the startup parameters. If specified for an RSR-capable active IMS, the specification of this parameter must match in the procedures used at the active site and at the remote site (after remote takeover). If XRF with MNPS is used, the USERVAR parameter is ignored.

For an XRF active IMS, this parameter overrides the XRF USERVAR specified in the IMS procedure library member DFSHSBxx.

- MTOID=*n* (in the IMS procedure)

This parameter must be part of the tracking DFSRSRxx procedure library member. Which DFSRSRxx member you use depends on the specification of the RSRMBR parameter in the tracking IMS procedure. It refers to which MTO name is to be selected from the TERMINAL macro's NAME list for the tracking IMS. This number (1, 2, or 3) also indicates which password is to be selected from the COMM macro's PSSWD list when opening the VTAM ACB for the tracking IMS. As many as three names can be specified in an XRF environment or an RSR tracking environment. You cannot override this parameter (as there is for APPLID) in the startup parameters. The default value is MTOID=1; only for a tracking IMS can you specify MTOID=3.

**Related reading:** For more information on these parameters, see *IMS Version 11 System Definition*.

The HSBID= parameter is ignored for a tracking IMS; the information is obtained instead from the MTOID parameter in the DFSRSRxx member and from the APPLID3 parameter in the IMS procedure. Whether or not your active IMS uses XRF, the tracking IMS always uses NODE3 (from the TERMINAL macro), PSSWD3 (from the COMM macro), and APPLID3 (from the IMS procedure).

The TERMINAL macro has three NAME parameters for VTAM terminals. Note that *name3* is used only for a tracking IMS. The COMM macro allows you to specify a third APPLID and password for the tracking IMS.

The following table shows the possible combinations of system definition and procedure parameters for RSR.

*Table 53. Relationship between IMS DC system definition and IMS procedures for RSR.*

| Terminal parameter | HSBID=1           | HSBID=2           | MTOID=3                             |
|--------------------|-------------------|-------------------|-------------------------------------|
| APPLID             | Application Name1 | Application Name2 | Set by APPLID3 or system definition |
| PASSWD             | Password1         | Password2         | Password3                           |

*Table 53. Relationship between IMS DC system definition and IMS procedures for RSR (continued).*

| Terminal parameter                   | HSBID=1 | HSBID=2 | MTOID=3 |
|--------------------------------------|---------|---------|---------|
| NAME (for master terminal)           | Name1   | Name2   | Name3   |
| NAME (for secondary master terminal) | Name1   | Name2   | Name3   |

**Notes:**

- For the master terminal, Name1 cannot be the same as Name2. The same is true for the secondary master terminal.
- Name3 for both the master terminal and the secondary master terminal can match the MTO names in the HSBID1 or HSBID2 positional parameter of NAME. In this case, all master terminals involved must be defined as VTAM local terminals.
- You need to specify MTOID=3 for the tracking MTO when you have an active and alternate XRF IMS at the active site, you want a different terminal for the tracking IMS, and you want to use the same system definition that the active IMS is using. This parameter is only valid when used for an RSR tracking IMS.

See *IMS Version 11 System Definition* for more information.

The following figure shows a possible configuration for a single IMS instance with XRF, RSR, and XRF after remote takeover. In the figure, it is assumed that separate master and secondary master terminals exist for each of the four operating systems that the active, alternate, or tracking IMS systems can run on. At any given time, however, only one active, one alternate, and one tracking IMS is started.

The figure also shows a master terminal configuration. RSR supports up to four separate Master and Secondary Master Terminals from a single system definition.




Figure 82. Master terminal configuration

Table 54 on page 678 presents another way of viewing the system definition and procedure parameters and how they interrelate.

The parameters and their sources are:

| Parameter     | Source                  |
|---------------|-------------------------|
| MTOID=n       | DFSRSRxx PROCLIB member |
| PSSWD=(a,b,c) | COMM macro              |

**APPLID=(a,b,c)**  
 COMM macro

**APPLID1,2,3**  
 Startup procedures

**USERVAR= or MNPS=a**  
 Startup procedures

*Table 54. IMS, XRF, and RSR configurations. How They Relate to Various System Definition and Procedure Parameters*

| Parameter                    | Basic IMS | Basic IMS with XRF | Basic IMS with RSR | IMS with RSR; XRF at Active Site | IMS with RSR; XRF at Remote Site | IMS with RSR; XRF at Both Sites |
|------------------------------|-----------|--------------------|--------------------|----------------------------------|----------------------------------|---------------------------------|
| MTOID=n                      | N/A       | N/A                | Y                  | Y                                | Y                                | Y                               |
| PSSWD=(a,b,c)                | a         | a,b                | a,b,c              | a,b,c                            | a,b,c                            | a,b,c                           |
| APPLID=(a,b,c)               | a         | a,b                | a,b,c              | a,b,c                            | a,b,c                            | a,b,c                           |
| APPLID1=, APPLID2=, APPLID3= | N/R       | N/R                | a                  | a,b                              | a,b                              | a,b                             |
| USERVAR= or MNPS=            | N/A       | N/R                | N/R                | Y                                | Y                                | N/R                             |
| XRF USERVAR or MNPS ACB name | N/A       | Y                  | N/A                | Y                                | Y                                | Y                               |

**Notes:**

1. Symbols:  
 N/A Not Applicable  
 N/R Not Required
2. For the APPLID3 keyword in the startup procedures, the “c” APPLID name, for the tracking IMS, is always optional and only serves as an override. In cases where three names are possible, all three names can be specified, and serve as overrides if specified on the startup procedures.
3. All systems are assumed to be running with VTAM terminals. The XRF configuration requires that the VTAM application name be the same as the ACB name.  
**Recommendation:** Make the VTAM application name the same as the ACB name for all configurations.

The following list provides detail on each configuration used in the table:

**Basic IMS**

The simplest IMS system. Nothing new is required for an IMS in this configuration. However, the APPLID in the startup procedures allows the APPLID in the COMM macro to be overridden. Note that the PSSWD in the COMM macro cannot be overridden. PSSWD selection for the tracking IMS is controlled by the MTOID parameter.

**Basic IMS with XRF**

The simplest XRF complex. Nothing new is introduced, except the startup parameter override option for the APPLIDs in the COMM macro. The USERVAR or MNPS startup parameter is not required. If specified, it overrides the XRF USERVAR or MNPS ACB name specified in the DFSHSBxx PROCLIB member.

**Basic IMS with RSR**

The simplest RSR configuration. A single IMS is started when the tracking IMS detects takeover for the single active IMS. The MTOID parameter is required by the tracking IMS to select the proper MTOs from the primary and secondary MTO-names lists. APPLID parameters are required in the startup parameters, and they must be different in the IMS procedures used

at the active and remote sites. Similarly, the USERVAR or MNPS parameters must match in the IMS procedures at the two sites.

#### **IMS with RSR and XRF at Active Site**

In this configuration, the XRF complex exists only at the active site. A single IMS will be started at the remote site.

The MTOID parameter is required by the tracking IMS to properly select the MTOs from the primary and secondary MTO-names lists.

The APPLID parameters are required in the startup procedures and they must be different in the procedures that are used at the active and remote sites. The APPLIDs on the startup procedures must be carefully specified: APPLID1 at the active site must not match the APPLID1 at the remote site. That is, the active site should specify  $APPLID1=a$ , and the remote site should specify  $APPLID1=b$ .

The USERVAR or MNPS startup parameter is required in the IMS procedure that will be used to bring up the active IMS at the remote site after a remote takeover. The USERVAR or MNPS startup parameter should not match the APPLID parameters in the startup procedures at this site. The USERVAR or MNPS startup parameter at the tracking site must match the USERVAR or MNPS ACB name that is being used at the active XRF site.

#### **IMS with RSR and XRF at Remote Site**

In this configuration, the XRF complex exists only at the remote site.

The MTOID parameter is required by the tracking IMS to select the proper MTOs from the primary and secondary MTO-names lists.

The APPLID parameters are required in the startup procedures, and they must be different in the procedures used at the active and remote sites. The APPLID parameters on the startup procedures must be carefully specified: APPLID1 at the active site must not match the APPLID1 at the remote site. That is, the active site should specify  $APPLID1=a$ , and the remote site should specify  $APPLID1=b$ .

The USERVAR or MNPS startup parameter in the RSR procedure is required in the startup procedure used at the active, non-XRF, IMS site. This parameter should not match the APPLID in the startup procedures at this site. This parameter must match the XRF USERVAR or MNPS ACB name that is used at the remote site after a remote takeover.

#### **IMS with RSR and XRF at Both Sites**

In this configuration, the XRF complex exists at both the active and RSR sites.

The MTOID parameter is required by the tracking IMS to select the proper MTOs from the primary and secondary MTO-names lists.

The APPLID parameters are required in the startup procedures and they must be different in the procedures used at the active and remote sites. The  $APPLID=(a,b)$  on the startup parameters must be carefully specified: All APPLIDs must be unique. That is, the active site should specify  $APPLID=(a,b)$ , and the remote site should specify  $APPLID=(c,d)$ .

The USERVAR or MNPS parameter is not required at either site. If specified at either site, it overrides the XRF USERVAR or MNPS ACB name, and must also be specified at the other site.

---

## XRF and RSR

If you consider the recovery support offered by XRF and by RSR, you might conclude that RSR function is a superset of XRF function. However, important differences exist between the kind of recovery support provided by XRF and that provided by RSR.

The main difference is the goal of each: XRF is intended to facilitate recovery in cases of brief or repairable failure of the active IMS, whereas RSR is intended to facilitate recovery in cases of lengthy or catastrophic failure of the active IMS. Several other important differences are listed in the following table.

*Table 55. Comparison of recovery functions provided by XRF and RSR*

| XRF                                                                                                                                       | RSR                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Relies on the same physical databases and many of the same physical resources as the active IMS, thus yielding a single point of failure. | Duplicates all physical resources without a distance limitation, thus no single point of failure exists, and so an area-wide physical problem can be overcome.                                                                                                                                                            |
| Supports DB/DC and DCCTL.                                                                                                                 | Supports DB/DC, DBCTL, DCCTL, and batch DL/I.                                                                                                                                                                                                                                                                             |
| Performs takeover on an IMS-by-IMS basis.                                                                                                 | Remote takeover includes all IMS systems that share databases.                                                                                                                                                                                                                                                            |
| No exposures to marooned data.                                                                                                            | <ul style="list-style-type: none"><li>• For planned takeovers, RSR has no exposures to marooned data.</li><li>• For unplanned takeovers, RSR users have exposures to marooned data.</li></ul>                                                                                                                             |
| Switching to the alternate IMS and back again is relatively easy.                                                                         | <ul style="list-style-type: none"><li>• For planned takeovers, switching back to the original active site is more complex than for XRF takeovers.</li><li>• For unplanned takeovers, switching back to the original active site is very difficult, and requires a planned takeover by the original active site.</li></ul> |
| Requires some IMS system management.                                                                                                      | RSR requires more IMS system management because of the need to replicate descriptors, programs, and other resources and objects at a second site.                                                                                                                                                                         |
| Takeover is fast, for example, one minute.                                                                                                | Takeover is slower, for example, one hour.                                                                                                                                                                                                                                                                                |

You can use RSR and XRF together or separately. After completing an RSR takeover, you can start an XRF alternate IMS to support the new active IMS (at the remote site). No XRF-alternate support exists for an IMS running as a tracking IMS.

Conversation is established prior to XRF takeover to avoid delay in sending log data after the XRF takeover. If any log data is missing after the XRF takeover, the tracking IMS obtains it.

---

## Defining an RSR environment with XRF

If your installation already has an XRF implementation, defining the environment to include RSR is not difficult. Defining TMS APPLIDs in an XRF complex is similar to defining them in a multiple-CPC data sharing environment. The TMS startup commands differ from site to site only by their APPLIDs.

See “Defining an XRF/RSR environment: active IMS and active site,” “Defining an XRF/RSR environment: alternate IMS and active site,” and “Defining an XRF/RSR environment: tracking IMS and remote site” for examples of defining RSR in your environment.

### Defining an XRF/RSR environment: active IMS and active site

```
SET APPLID(TMSA) ...
DEFINE SYSTEM(TMST,TMSX)
START TMS
START SYSTEM(ALL)
```

### Defining an XRF/RSR environment: alternate IMS and active site

```
SET APPLID(TMSX) ...
DEFINE SYSTEM(TMSA,TMST)
START TMS
START SYSTEM(ALL)
```

### Defining an XRF/RSR environment: tracking IMS and remote site

```
SET APPLID(TMST) ...
DEFINE SYSTEM(TMSA,TMSX)
START TMS
START SYSTEM(ALL)
```

Be aware that setting a different instance name for each TMS requires the active and alternate IMS systems to use different DFSRSRxx members, where TMINAME(xxxx) refers to the appropriate instance name. If online change MODSTAT tracking is used, both RSR members must have the same TRKMODS definitions.

**Recommendation:** Make all TMSs have the same instance name to allow the active and alternate IMS systems to use the same DFSRSRxx member. (Defining different instance names is only useful when multiple TMSs are running on the same CPC.)

The tracking IMS uses the SSIDs and RSENAME for tracking the XRF active and alternate IMS systems. When an XRF takeover occurs, the tracking IMS is aware that the log data that the new active IMS sends is a continuation of the log data being received from the old active IMS.

For ILS in a multiple-CPC environment, you can choose to start ILS on any one or any number of CPCs:

#### Starting ILS only on the alternate system's CPC

Because the workload is generally less on the alternate system's CPC, starting ILS only on the alternate system's CPC can be beneficial. However, doing so creates a situation in which, during an XRF takeover, both ILS and takeover processing must occur simultaneously.

#### Starting ILS on multiple CPCs

If the ILS that is engaged in conversation with the remote site fails, the remote site can obtain a conversation with another ILS.



### Starting ILS on less than all CPCs.

Choosing the CPCs on which to start ILS depends on your installation's requirements.

---

## Data sharing and RSR

RSR supports data sharing for IMS databases. All active IMS systems that share data send log data to a single tracking IMS. The tracking IMS is responsible for tracking all database activity for all of the active IMS systems in the service group.

When the remote site receives log data, it records it on SLDS and notifies DBRC of the data's existence. If the remote site is operating at the database readiness level, log data from all active IMS systems is passed to the database tracking IMS systems. RSR ensures that data integrity is maintained at the remote site for logs received both from active IMS systems participating in block-level data sharing and from non-data-sharing active IMS systems.

When a remote takeover occurs, all IMS systems in the active service group are affected. If you initiate a remote takeover because one IMS system in a data sharing environment fails, the workload of the entire service group is transferred to the remote site. Therefore, using RSR to provide XRF-like coverage for isolated disruptions is probably inappropriate.

The following figure shows an RSR data-sharing environment. Active IMS A and active IMS B share master database X and master database Y. The logs from active IMS A and active IMS B are sent to the single tracking IMS at the remote site. At the remote site, the tracking IMS has shadow database X, shadow database Y, SLDS A, SLDS B, and its own log.




Figure 83. Data sharing in an RSR complex. The log data for the active IMS systems is sent to the tracking IMS.

The tracking IMS manages data sharing so that it:

- Avoids physical sharing of databases at the tracking IMS, thus avoiding the complexity and cost of locking and data sharing at the remote site
- Logically merges all log records from the data sharing subsystems, thus preventing transaction and database application inconsistency
- Uses a single tracking IMS to support many active IMS systems

## DRD and RSR

Dynamic resource definition (DRD) is supported on the IMS systems at the active site. DRD is not supported on the Remote Site Recovery (RSR) tracking system. When a remote takeover occurs, DRD is supported on the new active IMS systems.

IMS active site processing with DRD is tracked by the RSR tracking system. However, DRD is not used by the RSR tracking system. Once remote takeover has occurred, DRD is used by the new active IMS systems.

If DRD is enabled on the IMS systems at the active site, type-2 commands can be used to dynamically create, update, or delete database, program, routing codes, and transaction resources and descriptors, without requiring a cold start or a MODBLKS online change. Any changes that are made dynamically are logged with an X'22' log record and are recoverable across a warm or emergency restart. In order to recover changes across a cold start (or perform an emergency restart with the COLDSYS parameter), save the resource and descriptor definitions of your IMS to a system resource definition data set (RDDS), so that they can be imported from the system RDDS during cold start. The automatic export function or the EXPORT command can be used to save resource and descriptor definitions to an RDDS.

Each IMS at the active site with automatic import or export enabled must have its own set of system RDDSs defined. System RDDSs cannot be shared between IMS systems. System RDDSs contain an IMSID identifier that is checked for validity during IMS initialization. Each IMS at the remote site that is brought up after remote takeover with automatic import or export enabled must have access to its own set of system RDDSs. The RDDSs at the remote site can either be the same physical data sets as those used by the old active, or identical images of the RDDSs at the active site.

Definitional changes (using DRD) made to resources at the active site are not automatically made to the resources at the tracking site. It is up to the user to perform a MODBLKS online change to make similar changes to the resources on the tracking IMS.

The X'4004', X'4006', X'4007', X'4083', and X'4012' checkpoint log records must be transported from the active site to the tracking site to enable the remote IMS systems to warm start successfully during a remote takeover.

In order for the IMS systems at the tracking site to be able to successfully perform an emergency restart, the X'22' resource change log records are also needed. Do not filter out the X'4004', X'4006', X'4007', X'4083', and X'4012' checkpoint log records and the X'22' records using the DFSFTFX0 exit routine. They must be transported from the active site to the remote site.

### **Restarting the active IMS systems after a remote takeover (planned and unplanned)**

If the new active IMS systems are warm or emergency restarted, the runtime resource definitions for database, program, routing code, and transaction resources and descriptors are rebuilt from the log records.

If the new active IMS systems are cold started with DRD enabled and automatic import enabled, the resource definitions for program, database, routing code, and transaction resources and descriptors are imported from the system RDDS with the most current data. In order for the new active IMS to come up with the same resource definitions as the IMS at the old active site, a copy of the old active IMS system's most current RDDS must exist at the remote site.

The most current RDDS from the active site must either be made available to the IMS taking over at the remote site or a copy of it must be built at the remote site. This can be achieved by the use of data set mirroring and the XRC tracking

function of RSR or by using the DRD utilities that are provided with IMS and are accessible through the IMS Application Menu. For example, the DFSURCL0 utility can be used to recreate the most current RDDS from the X'22', X'4001, X'4004', X'4006', X'4007', X'4083', X'4098', and X'45FF' log records that are transported from the active system to the remote site, or the DFSURCP0 utility can be used to copy the contents of an existing RDDS into a new RDDS.

---

## Tracking an IMSplex

A tracking IMS can connect to a Common Services Layer (CSL) to enable automatic RECON loss notification or to receive commands from the Operations Manager (OM). If the active site is an IMSplex with CSL, then the remote site must be a different IMSplex with CSL.

---

## RSR log management

The following topics describe RSR log management for the active and tracking IMS systems.

### Active IMS

The active IMS uses a control region, DL/I separate address space (DLISAS) region, DBRC region, dependent regions, and optionally an IRLM region. RSR also adds a new subsystem for the Transport Manager Subsystem (TMS).

The IMS logger establishes a conversation with the log router component of the tracking IMS using the TMS of the active IMS. As log buffers fill, they are sent to the tracking IMS, prior to OLDS I/Os being initiated. When it is at the tracking IMS, the data is not routed until the tracking IMS confirms that the data has been written to DASD at the active IMS.

When disruptions occur and log data is produced at the active site but not sent to the tracking IMS, the tracking IMS sends a request to the isolated log sender component at the active site to send all log data created during the disruption. This call requests all log data created during the disruption. The isolated log sender obtains data set information from DBRC and then reads and sends the data to the log router of the tracking IMS.

### Tracking IMS

The tracking IMS operates rather differently from the active IMS. The tracking IMS uses a control region, DBRC region, and a region for the TMS. The tracking IMS also uses a DLISAS region unless you specify TRACK=RLT or LSO=Y, both of which tell the tracking IMS not to use a DLISAS region; this can save you the additional CPC and storage resource of maintaining the DLISAS region. The DLISAS region is only required for DL/I databases; it is not required for Fast-Path-only systems.

The log router component runs in the control region, acting as the source of log data for other components. The log router establishes a conversation with the ILS and the loggers of all operational online active IMS systems in the active service group.

As log data is received from active IMS systems, the log router writes the data on an SLDS. This SLDS is different than the usual IMS SLDS in that it is not an archive of an OLDS nor is it a log from a batch job; it is unique to the tracking IMS in an RSR complex. The tracking IMS informs DBRC of SLDS data set creation and all the log-data-set-to-database relationships. SLDSs are recorded in PRILOG and PRISLDS records in the tracking RECON data set.

The tracking IMS uses its own log data sets to record information required for tracking IMS restart and recovery, such as information about SLDS data set names and current positions in the data sets. DBRC tracks the subsystem log in the RECON data set in the same way as it does in the active IMS. Tracking IMS log data sets are recorded in PRIOLDS and PRITSLDS records in the tracking RECON data set.

After a network or tracking IMS disruption, gaps undoubtedly exist in the log data. After such a disruption, the log router reestablishes a conversation with the active IMS logger and begins receiving log data again. When the log router detects the gap in the log data, it sends a request to the isolated log sender to send the missing log data. The log router then begins routing log data to the database tracking IMS systems, obtaining the log data either from the local SLDS or the active IMS.

## Example of an RSR complex

The following figure illustrates some of the concepts introduced in this topic, and shows what an RSR complex with XRF and data sharing might look like.




Figure 84. Example of an RSR complex with XRF and data sharing

The IMS logs are sent from the active IMS systems as they are created. This RSR complex uses XRF and data sharing. The active site has two active IMS systems, A and B, which are configured with XRF and data sharing. Each IMS has its own logs, but both IMS systems share a master database, a RECON data set, and a

coupling facility, which includes message queues and other structures. The remote site contains one tracking IMS that is used for storing log data received from the active IMS systems. All IMS systems in the RSR complex are configured with Transport Manager Subsystems (TMSs).

If one of the active site IMS systems is disrupted, you should use its XRF alternate IMS to take over its processing. If both IMS systems at the active site are disrupted, this setup is designed on the assumption that you are willing to recover both IMS systems, A and B, on the tracking IMS.

When updates are made to an active IMS, this IMS writes log data to an online log data set (OLDS), or, for batch, to a system log data set (SLDS). At the same time as the disk write, the active IMS also sends this log data to the tracking IMS. The tracking IMS stores this data in SLDS.

At the remote site, DBRC maintains log and database recovery information for the tracking IMS. The RECON data set of the tracking IMS is not a mirror of the RECON data set at the active site. DBRC records log data received from both active IMS systems and maintains database recovery information for the tracking IMS in the RECON data set of the tracking IMS.

RSR does not support active application processing against the shadow databases until a remote takeover of active processing occurs. At that point, the remote site becomes the new active site and you can restart the active IMS systems at that site. If the active IMS participates in data sharing, you need to switch all sharing IMS systems (the entire service group) to the remote site on an RSR takeover.

---

## Coordinated IMS and DB2 for z/OS recovery support

In the event you have an unplanned or planned takeover, and you have both IMS and DB2 for z/OS databases, you can simplify your data recovery of both systems by coordinating RSR with 3990 Extended Remote Copy (XRC).

In DB2 for z/OS data-sharing environments where XRC transmits DB2 for z/OS data to a remote site for recovery, RSR support is provided. RSR with XRC allows IMS and DB2 for z/OS logs to be synchronized and ensures that IMS and DB2 for z/OS databases are consistent when a planned or unplanned takeover occurs.

The following list discusses how XRC supports RSR and lists the requirements for XRC.

- How XRC supports RSR for IMS and DB2 for z/OS

The following figure illustrates the concept of how an XRC communicates with an RSR Complex with data-sharing. The DB2 for z/OS data-sharing environment is also split into an active site and a remote site. The active site contains the DB2 for z/OS subsystem, DB2 for z/OS logs, boot-strap data sets (BSDS), and an XRC component. The remote site contains another XRC and copies of the active site's DB2 for z/OS logs and BSDS. The tracking IMS at the remote RSR site for IMS communicates with XRC at the remote site for DB2 for z/OS.




Figure 85. Example of an RSR complex with XRC tracking

To support the synchronization of the IMS and DB2 for z/OS logs, the IMS shadow logs run by RSR are always kept behind the DB2 for z/OS shadow logs, which are run by XRC. The logs are synchronized when the user captures the IMS RSR truncation time stamp and provides it as the ENDLRSN in the conditional restart records for the DB2 for z/OS restarts. Because IMS shadow logs are always kept behind DB2 for z/OS shadow logs, the IMS tracking log end point will always be earlier than the end point of the DB2 for z/OS logs. When the DB2 for z/OS systems are restarted at the remote site after takeover, the DB2 for z/OS logs will be truncated based on the IMS truncation time stamp. Because the IMS and DB2 for z/OS tracking IMS systems end at a consistent point, the IMS and DB2 for z/OS logs are synchronized.



The following figure shows the IMS and DB2 for z/OS log transmission to the remote site. The IMS and DB2 for z/OS log transmissions to the remote site start at different times. When the remote site takes over, RSR chooses a point in time when both the changes are routed or committed (timestamp t1) in one of the IMS logs (IMS B) and when the other IMS and DB2 for z/OS subsystems are still viable. RSR truncates the other logs to time t1 for start purposes so that all of the logs have the same start time.




Figure 86. IMS and DB2 log transmission to a remote site

To use the coordinated IMS and DB2 for z/OS recovery support, the XRC function must be used for DB2 for z/OS logs and bootstrap data sets (BSDSs) only. The tracking IMS that supports XRC can track active IMS systems that do not support XRC. XRC is enabled when you specify an XRC session ID as a startup parameter (in the DFSRSRxx PROCLIB member) for the tracking IMS. The volumes containing the active logs and BSDSs for the DB2 for z/OS subsystems that need to be synchronized with the IMS systems in the RSR complex must be shadowed by this XRC session.

- Requirements for RSR and XRC tracking

The requirements for coordinated IMS and DB2 for z/OS recovery support are:

- In environments where the production workload is spread across multiple CPCs, RSR requires a 9037 sysplex timer (or equivalent). XRC also requires a sysplex timer for the application programs in multi-CPC environments.
- The tracking IMS and the system data mover (SDM) (part of z/OS that interacts with the storage controls that control the XRC primary volumes for

the session) must be on the same operating system in order for the tracking IMS to coordinate processing with an XRC session.

- The tracking IMS must have RACF authorization to make XRC requests through the SDM SPI. You must authorize the user ID associated with the tracking IMS control region to have read access to the FACILITY class profile of STGADMIN.ANT.XRC.COMMANDS.
- The XRC session must only have DB2 for z/OS logs and BSDSs and not DB2 for z/OS databases.

---

## Installing RSR

One of the most important things you must consider before you install an RSR complex is the replication of hardware and software at the two (or more) sites. Then you must consider the configuration of the active site and, more importantly, the remote site, including the databases and data areas.

### Hardware replication

In general, hardware need not be replicated at the remote site.

The tracking CPC must have the same architecture and hardware features as the active IMS, unless the use of these features by IMS is optional or is controlled by specifications other than IMS system definition.

The remote site CPU should be large enough to support the expected workload after a remote takeover. However, the tracking CPC need not have the same capacity as the active CPU, because a small tracking CPC can support a large active CPC if the readiness level of recovery (RLT) is used. For either recovery (RLT) or database (DLT) readiness levels, the active site CPC and the remote site CPC can be different sizes.

Your external storage devices (DASD and tape units) need not be the same at the active and remote sites. The tracking DASD must have the same architecture and hardware features as those at the active site, unless the use of these features by IMS is optional or is controlled by specifications other than IMS system definition.

If you use DASD with different track capacities than those at the active site, you should use block sizes at the remote site that match the lowest block size at the active site, otherwise you must manage the additional differences, such as for allocation specifications. For database data sets, the VSAM CI size and the OSAM block size must be the same at both active and remote sites.

**Recommendation:** If you use Fast Path and sequential dependents (SDEPs), use the same device types for your active and tracking DEDB areas. The last usable CI varies by device type, and using the same type for your DEDB areas ensures that the tracking IMS can track SDEP inserts when you reach the end of the area.

**Related reading:** For more information on SDEPs, see *IMS Version 11 Database Administration*.

### Software replication

Because the remote site's IMS is effectively a duplicate of the active site's IMS, any changes to system definitions, options, application programs, and IMS maintenance levels generally need to be made to both IMS systems. You are responsible for

replication of code and definitions (such as PSB and DBD definitions), and you must replicate any changes made to them.

As is the case with an XRF alternate IMS, the RSR tracking IMS must use the same release level of IMS (but not necessarily the same maintenance level) as does the active IMS. You should also use the same z/OS release level for each IMS, but at a minimum, both sites must have the lowest level of z/OS required by the current release of IMS running at the sites. Normally, the IMS and z/OS systems should be similar enough at each site so that the output of the stage 2 of IMS system definition can be copied from one IMS to another.

Certain requirements exist for the placement and replication of data sets for RSR. Specifically:

1. Data sets whose content must be replicated at remote sites

These are data sets containing definitions, procedures and code (ACBLIBx, DBDLIB, FORMATx, JOBS, MODBLKSx, MODSTAT, PGMLIB, PROCLIB, PSBLIB, SDFSRESL, and TFORMATx). With the exception of the MODSTAT data set, you are responsible for replicating any changes to these data sets. Changes at the active site to the MODSTAT data set are tracked at the remote site.

If your remote site is tracking more than one active site, its resources are likely to be a composite of all the IMS systems it tracks. In this case, it might be necessary to have separate tracking data sets for the ACBLIBx, DBDLIB, and PSBLIB data sets. You can build a proper ACBLIBx by merging the DBDLIBs into a tracking IMS DBDLIB, merging the PSBLIBs into a tracking IMS PSBLIB, and then generating ACBs for the tracking IMS. The tracking IMS can use any FMTLIBx that contains the required formats; you do not need to make a composite FMTLIBx.

2. Data sets for which space must be allocated at the remote site and whose content must be created specifically for the site

These are the RECONx data sets. The contents of the RECON data sets are created by DBRC commands and by the tracking IMS as it tracks the active IMS systems.

The RECON data set at the remote site is neither a physical nor logical copy of the RECON data set at the active site.

3. Data sets for which space must be allocated at remote sites

These are all other IMS data sets not already mentioned, including the data sets that are shared by an XRF active and alternate IMS systems.

Database data sets need not have space allocated if they are not tracked or if they are not going to be maintained at the database readiness level.

Data set names do not need to be the same at the different sites. When data set names are different, these differences must be reflected in your JCL procedures, DBRC definitions, and dynamic allocation descriptors. Subsequent alterations to these procedures, definitions, and descriptors must then be made manually at each site as appropriate.

Any alterations made to z/OS or to its supporting product definitions and data sets, in order to install or tune IMS, must normally be replicated at remote sites.

## **DL/I databases**

In preparing for installation of an RSR complex, you need to consider the following for your DL/I databases:

### Database coverage

RSR supports the following DL/I database organizations: HDAM, PHDAM, HIDAM, PHIDAM, HISAM, SHISAM, and PSINDEX. You can decide to track all databases managed by a given active IMS or a subset of the databases. All databases that are to be tracked must be registered with DBRC.

If a database is not to be tracked (that is, not covered), RSR does nothing for that database: RSR does not maintain recovery data in the remote site RECON data set for this database, nor does it maintain a shadow database.

If only a subset of the databases managed by an active IMS are to be tracked, you should consider all associated database relationships, such as logical relationships or primary or secondary index relationships, and select coverage for these related databases accordingly. You also need to consider databases used by certain critical applications and select coverage for those groups of databases accordingly.

RSR does not automatically track databases that have logical relationships or primary or secondary index relationships with other databases, nor does it track databases used by an application that has all its other databases tracked; RSR only tracks those databases you ask it to track.

If you do not consider these other relationships, a database after a remote takeover might be physically available and current but be logically unavailable, because processing of the database requires the availability of related non-tracked databases.

### Installation of shadow databases

For all those databases that are to be tracked, at either recovery (RCVTRACK) or database (DBTRACK) readiness level, you need to send an image copy of each database to the remote site. Then you must tell DBRC about each database, using the NOTIFY.IC command, and register the databases, using the INIT.DB or INIT.DBDS command. The image copy is used as the base to recover the shadow database (for recovery readiness level) or as the base to begin tracking the database (for database readiness level).

For recovery readiness level tracking (RCVTRACK), you can store the database image copies on any storage medium. After a remote takeover, the databases must be forward recovered before they are ready for updates.

**Recommendation:** Run online forward recovery (OFR) for all databases before a remote takeover.

For database readiness level tracking (DBTRACK), the image copies must be applied to the database data sets in order to be ready for updating during tracking. You install the shadow database using the GENJCL.RECEIVE command. DBTRACK on an RLT system is treated the same as for RCVTRACK databases.

### Generating DBDs

The database definitions (DBDs) at each site must match. Database data set DD names must also match, because they identify the data sets.

### Generating PSBs

The program specification blocks (PSBs) at each site must match. PSBs are not needed at the remote site until immediately before restarting the tracking IMS as the new active IMS.

### **Generating ACBs**

The generation of application control blocks (ACBs) at each site must be coordinated so that the members are kept identical. You can run the ACB Maintenance utility (DFSUACB0) at the remote site (rather than copying the ACB library from the active IMS system), but you must supply matching DBD and PSB input. If the tracking IMS is tracking more than one active IMS, you must perform a composite ACB generation.

### **System definition**

System definition for a database readiness level tracking IMS must include all databases that are to be tracked.

If only one active IMS is to be tracked, that system definition of the active IMS system can be used for the tracking IMS. If more than one active IMS is to be tracked and none of the definitions from the active IMS system includes all the databases to be tracked, you must perform a special tracking IMS system definition.

### **Dynamic allocation**

Database data set names need not be the same at the active and remote sites. You can choose to assign different data set names for asset identification and control purposes.

## **Fast Path databases**

In preparing for installation of an RSR complex, you need to consider the following for your Fast Path databases:

### **Database coverage**

RSR supports DEDB areas, including VSO DEDBs; MSDBs are not supported. You can decide to track all databases and areas managed by a given active IMS or a subset of the databases and areas. All databases and areas that are to be tracked must be registered with DBRC.

### **Multiple area data set (MADS)**

RSR supports multiple area data sets (MADSs) for the shadow copies of the database areas. The number of area data sets (ADSs) for each area at the remote site can be different from the number at the active site. Because no dependent regions are in the tracking IMS, the ADS Create utility and ADS Compare utility cannot be used at the remote site. ADS DSNAMES and DD names need not be the same at the active and remote sites.

You should create additional ADSs by copying (for example, using the z/OS Access Methods Services REPRO command) the ADS just installed, and then make the new ones available using the `CHANGE.ADS AVAIL` command.

### **Installing shadow databases and areas**

For all those areas that are to be tracked, at either recovery (RCVTRACK) or database (DBTRACK) readiness level, you need to send an image copy of each area to the remote site. Ensure that your databases and areas are registered in RECON. Send image copies to the remote site and register these copies in RECON. The image copy is used as the base to recover the shadow database (for recovery readiness level) or as the base to begin tracking the database (for database readiness level).

For each covered area, when you format the area (at the active site), you must include the GSG name.

For recovery readiness level tracking (RLT), you can store the area image copies on any storage medium. After a remote takeover, the areas must be forward recovered before they are ready for updates. For database readiness level tracking (DLT), the image copies must be applied to the database in order to be ready for updating during tracking, thus eliminating the need to forward recover the areas after a remote takeover. You install the image copy using the GENJCL.RECEIVE command.

#### **Main storage databases (MSDBs)**

Because MSDBs are not supported at the remote site, MSDB-specific data sets, including MSDBCPn, MSDBDUMP, and MSDBINIT, are not required. If you send them to the remote site, these data sets are not tracked.

#### **Online change**

One MODSTAT data set must be defined for the tracking IMS, and one additional MODSTAT data set must be defined for each MODSTAT data set that is to be tracked. Thus, tracking a non-XRF-capable IMS requires one MODSTAT data set, whereas tracking an XRF-capable IMS requires two MODSTAT data sets. But for any given active IMS, only the MODSTAT information for the *currently active* IMS is tracked. The dynamic allocation member (DFSMDA) for the tracking data sets is defined in the DFSRSRxx PROCLIB member at the active site.

A tracking IMS that was previously an active IMS is not required to use the same DD names specified for it when it was an active IMS.

### **Running IMS workload on multiple z/OS images in an RSR environment**

In general, if you are running your IMS workload on more than one z/OS image in an RSR environment, you must use z/OS Global Resource Sharing (GRS) on each of the z/OS images.

This is because the isolated log sender (ILS) needs access to the IMS log data (OLDS or SLDS) on each of the z/OS images in the RSR environment. OLDSS in this environment should have unique data set names and exist on disks shared among the z/OS images.

#### **ILS and the OLDS on IMS TM or IMS DB**

Because the isolated log sender (ILS) can access data on disk without operator intervention, it is highly desirable to provide an active IMS with a large amount of OLDS space. With lots of space, isolated log data can be accessed and sent quickly in many cases, without the need for tape mounts to read archived SLDS data sets.

Isolated log data can be obtained from OLDS data sets only at specific points in the OLDS cycle. Because the ILS cannot access data on an OLDS until the active IMS is finished writing to it, each individual OLDS should not be made too large; this can delay the sending of isolated log data. IMS commands can be used to cause an OLDS switch, but this requires operator intervention.

Because the ILS is unable to read OLDS data after completion of archive, it is operationally desirable to be able to defer archive operations after network or remote site outages (assuming an adequate number of OLDS data sets are defined to allow continued IMS operation).



If it is ever possible that the OLDS archive utility (DFSUARC0) can be executed on a z/OS image different than the z/OS image where the ILS is running, you must use z/OS Global Resource Sharing (GRS) on the z/OS images. This is because a combination of z/OS ENQ/DEQ and DBRC are used to prevent exposure to the ILS reading an OLDS that is being re-used by a TM or DBCTL IMS system. After an OLDS has been archived, it is available for re-use by the TM or DBCTL system and must be considered unavailable for the ILS.

## ILS and batch IMS systems

If you allocate your batch log data sets on disk, conflicts with the ILS must be avoided. This is particularly true when the log data sets are reused or scratched, even after they are copied by the archive utility. Because of potential conflicts, the ILS allocates SLDS data sets with disposition OLD to prevent sharing and inadvertent scratching. In order for this to be effective when reuse or scratching can be executed on a z/OS image different than the z/OS image where the ILS is running, you must use z/OS Global Resource Sharing (GRS) on the z/OS images.

After a batch log has been copied and DBRC RECON updated with the change, DBRC provides the new data set to the ILS if the log data is required. Situations can exist where the previous name is given to the ILS and that data set is not yet sent. If this is the case, you can use the DISPLAY ILS command to check for pending data set transfers.

---

## Initializing RSR

Initializing RSR involves initializing the active site and initializing the remote site.

### Initializing the active site

To initialize the active site, perform the following tasks.

1. Initialize the Transport Manager Subsystem (TMS).

You can initialize the TMS with a z/OS START command or by submitting a job.

You can issue TMS commands from a SYSIN data set at component start, or from a z/OS console using the z/OS MODIFY command.

**Related reading:** For more information on the TMS commands, see *IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands*.

A typical command sequence for starting the TMS is:

- a. SET

The SET command allows you to specify various parameters that generally stay in effect while the TMS task is running.

- b. DEFINE SYSTEM

Use the DEFINE command to define one or more TMSs in different CPCs to support IMS components for one or more RSR global service groups.

You can specify one or more DEFINE SYSTEM commands when starting a TMS.

- c. START TMS

Use the START TMS command to make TMS functions available to other address spaces.

For each command, you see error messages if command parsing errors or command sequence errors occur. Otherwise, you see a message indicating successful completion of the command.

The following example shows sample commands used to start the TMS.

```
SET APPLID(SYSN1) APPLCOUNT(10) PASSWORD(HOH0)
DEFINE SYSTEM(SYSN2,SYSN3)
START TMS
```

The START TMS command causes a VTAM access-method control block (ACB) for the TMS to be opened and logons to be enabled. The TMS then attempts to allocate conversations with any other TMSs it knows about (from the DEFINE SYSTEM command). Whenever conversations are allocated between TMSs, the TMSs exchange global service group and local service group directory information. These TMS conversations are not used by other RSR components.

Dynamic definition of global service groups, service groups, and remote TMSs is supported. When an RSR component identifies itself to a TMS, its name is distributed to other TMSs, thus supporting a dynamic distributed directory. When a component issues an allocate request to a different component, the TMS supporting the requester sends the request to the TMS that supports the requested component. If the requested instance and component are found, allocation proceeds; otherwise, it fails.

2. Initialize the IMS logger.

The logger obtains the service group names for the tracking service group from DBRC. It then attempts to identify to and allocate a conversation with the log router at the remote site. If errors occur, various messages are issued, but logger initialization continues. Errors do not prevent subsequent attempts to initiate these initialization processes for online IMS systems.

The OLDS block size should be no larger than 32708 bytes. If the OLDS block size is larger, then the logger does not establish a conversation with the tracking IMS, and real-time log transport is not available. The log data is sent to the remote site at a later time by the isolated log sender.

3. Initialize DL/I batch.

A batch job specifies the GSG name when it signs on to DBRC. Batch is not allowed for tracking service groups; therefore, if the local service group is not the active one, the signon request is rejected. Batch jobs that do not specify a GSG name cannot update databases that have an associated GSG name; if update is attempted, the authorization fails.

During initialization, the batch logger obtains the name of the remote service group from DBRC. The logger attempts to identify to the TMS and establish a conversation with the tracking IMS. If the logger cannot establish conversations with any tracking IMS systems, real-time log transport is not available for batch execution. The batch log data set is sent to the remote site at a later time by the isolated log sender.

The batch log data set block size should be no larger than 32708 bytes. If this block size is larger, then the logger does not establish a conversation with the tracking IMS.

## Initializing the remote site

To initialize the remote site, set the execution parameters, initialize the subsystem, and start database tracking.

To initialize the remote site, perform the following tasks:

1. Set execution parameters and configuration.

The TRACK parameter of the IMS procedure determines whether the control region being started performs RSR tracking.

**Related reading:** For more information on the TRACK parameter, see *IMS Version 11 System Definition*.



Use the RSR(NO) parameter in the DFSRSRxx PROCLIB member of the active site to disable RSR functions. Disabling RSR function by using the RSR(NO) parameter is only necessary if you specify the GSGNAME in the system definition. The RSR() parameter is only meaningful for active IMS systems and DL/I database tracking is not initialized if you specify LSO=Y in the IMS procedure for the tracking IMS.

2. Initialize the Transport Manager Subsystem (TMS).

The TMS at the remote site is initialized in the same way as at the active site.

**Related reading:** For more information on initializing the TMS at the active site, see Initialize the Transport Manager Subsystem.

3. Initialize DL/I database tracking.

DL/I database tracking is only activated in a database level tracking IMS; that is, a subsystem for which TRACK=DLT is specified in the IMS procedure. DL/I database tracking is initialized during tracking IMS initialization. It creates and initializes the DL/I database tracking data space.

A DL/I separate address space (SAS) region is required to process log records that keep the shadow full-function databases updated. Start this region as you would in any normal IMS system.

The PST parameter in the EXEC statement of the IMS PROCLIB specifies the maximum number of PSTs used for DL/I database tracking. The MAXREGN number you specify during system definition is used, if it is not overridden. The number should be approximately twice the number of DL/I PSTs used in all of the active IMS systems tracked by this tracking IMS. As a minimum, 2 PSTs are used.

4. Initialize Fast Path database tracking.

Before Fast Path database tracking can begin, you must:

- Specify a database readiness level tracking IMS (specify TRACK=DLT in the IMS procedure).
- Specify Fast Path during system definition (include an FPCTRL macro in your system definition).

During initialization, Fast Path database tracking:

- Loads required modules
- Creates and initializes its data space

Some processes normally done during initialization (in non-RSR subsystems) are skipped to reduce storage requirements and minimize initialization time for the tracking IMS.

---

## IMS error handling for RSR for the active site

The following components are involved in error handling for the active site: Transport Manager Subsystem (TMS), Isolated log sender, Online logger, and DL/I batch logger.

### Transport Manager Subsystem (TMS)

The TMS isolates IMS components from TMS failures and from VTAM failures. The general procedure is to take an SDUMP for diagnostic purposes, recover damaged structures if possible, and retry if possible. Retry might result in a return code's being given to the IMS component invoking the function.

Failures that occur in the TMS address space do not, in general, cause termination of IMS components. If the TMS cannot recover from a particular failure, IMS components are informed of the failure and can

continue to use the existing conversations. At this point, new conversations cannot be allocated (on the active processor) because of the failure. If new conversations must be allocated, a new TMS must be brought up.

Do the following for all IMS active IMS systems at the same site as the failed TMS to start new TMS conversations for these active IMS systems:

1. Issue a `/STOP SERVGRP` command to break the communications with the old TMS.
2. Issue a `/START SERVGRP` command to start communications with the new TMSs.

These steps are required if a need for the isolated log sender arises after a TMS address space failure. Note that the current conversations between the active IMS and the tracking IMS continue to run, but any new TMSs brought up do not know about these conversations.

If a transport-manager-to-transport-manager conversation fails, an immediate attempt is made to reallocate the conversation, in case an alternate path is available through the VTAM network. If that attempt fails, periodic attempts to reestablish a conversation are made based on a user-specified timer interval. VTAM errors on TMS conversations are identified by TMS error messages (prefix ELX) for diagnostic use.

#### **Isolated log sender**

In the event of allocation failures or permanent I/O errors on an active IMS log data set, the isolated log sender attempts to use a dual copy of the log data set, if one exists. If no dual copy is available, or if the attempt to use one fails, you need to recover the log. After the log data set is recovered, you can issue a `/START ISOLOG` command on the tracking IMS to re-initiate isolated log transport processing.

**Related reading:** For more information on the `/START ISOLOG` command, see *IMS Version 11 Commands, Volume 2: IMS Commands N-V*.

If log recovery cannot be performed, you must create new database image copies and transport and reinstall them at the remote site.

Abnormal termination of an ILS instance or of the copy of DBRC in the TMS address space does not terminate the TMS. You can recover from these kinds of failures by using a `START ILS` command.

It is recommended that you run the ILS on at least two different z/OS subsystems. Although the log router communicates with only one ILS at a time, having a second ILS available in case of failure of the first ILS allows for continuous ILS function. Having a second ILS can prevent problems associated with:

- ILS failure
- DBRC failure leading to ILS failure
- Partial network failures
- TMS, z/OS, VTAM, or CPC failures

If you do have a second ILS acting as a standby, both ILSs must have access to the OLDS and SLDS data sets. In order to control access to the OLDS and SLDS across multiple z/OS images, you need to use z/OS global resource sharing (GRS).

**Related reading:** For more information on z/OS GRS, see “Running IMS workload on multiple z/OS images in an RSR environment” on page 695.

### Online logger

The IMS online logger sends log data to the tracking IMS before initiating I/O to the active IMS OLDS. The Log Filter exit routine is given control as part of the send process (see *IMS Version 11 Exit Routines*), but if a failure in the exit routine occurs, the routine abends.

Other than a VTAM overload, failures in normal operation of the send process result in termination of the conversation involved. VTAM overloading causes the online logger to suspend sending log data to the tracking IMS, but the conversation is not terminated. This does not preclude subsequent attempts to reestablish the conversation after it has been terminated. Unless you issue a `/STOP SERVGRP` command, the IMS online logger attempts to resume suspended conversations and reconnect terminated conversations at each OLDS switch. If you want a faster reconnect attempt, the MTO can issue the `/START SERVGRP` command.

Attempts to reestablish failed conversations can be prevented by issuing the `/STOP SERVGRP` command. You might want to prevent these attempts if communications to the remote site or the tracking IMS are expected to be lost for an extended period of time.

### DL/I batch logger

The batch logger sends log data to the tracking IMS before initiating I/O to the active site SLDS. The Log Filter exit routine is given control as part of the send process (see *IMS Version 11 Exit Routines*), but if a failure in the exit routine occurs, the routine abends.

Failures in normal operation of the send process result in termination of the conversation involved.

---

## IMS error handling for RSR for the remote site

The following aspects are involved in error handling for the remote site: Log router, DL/I database tracking, Fast Path database tracking, online forward recovery, and online change.

### Log router

The following three general classes of error conditions are recognized by the log router:

#### System errors

System errors result from system resource shortages (for example, a shortage of storage) or are IMS internal logic errors.

Whenever possible, the effect of system errors is limited to the current "process". For example, if the log router is unable to obtain storage for the control blocks and data areas required to manage a conversation with an active IMS, the conversation is terminated. In this case, the active IMS log is obtained later during isolated log transport.

#### Communication errors

Communication errors affect the conversations between the log router and active site components (loggers and isolated log senders) or affect the connection between log router and the TMS.

Communication errors in a conversation result in the loss of contact with the active site component. The IMS operator is informed of the error by means of an error message. After the problem is resolved and communications are reestablished, normal

processing resumes. In the case of active loggers, the current log data is received from the logger, and any missing log data is obtained by the isolated log sender. If a conversation with the isolated log sender is lost, log transport resumes at the point of error when communications are reestablished.

If the connection between the log router and the TMS is lost, processing on current conversations continues, if possible. However, no new conversations can be established until the log router is reconnected with the TMS. After the problem is resolved, you must issue a `/STO SERVGRP` command followed by a `/STA SERVGRP` command to reconnect the log router to the TMS and all active IMS systems. You should wait until all batch jobs or BMPs are complete before issuing the `/STO SERVGRP` command.

If VTAM terminates, all conversations are lost. After VTAM and the TMS are restarted, you must issue a `/STA SERVGRP` command.

### Log media errors

Log media errors result from I/O errors on the tracked SLDSs. The following types of error can occur:

- Write error on a tracking SLDS

If single SLDS logging is used, or if an error occurs on both SLDSs, the current SLDS is closed and DBRC is notified that the SLDS is invalid. DBRC deletes the record of the invalid SLDS in the RECON data set. A message is issued identifying the SLDS, and a new SLDS is created; the log router proceeds, obtaining any missing log data.

If dual logging is used, the SLDS in error is closed immediately and the remaining SLDS is closed as soon as previously initiated writes have completed. DBRC is informed of the invalid SLDS. You can choose to ignore the problem and continue with a single copy of the SLDS, or you can copy the good SLDS to a new data set and use the DBRC `CHANGE.PRILOG` (or `CHANGE.SECLOG`) command to inform DBRC that the SLDS is valid.

- Read error on a tracking SLDS or archived SLDS

Read errors on tracking SLDSs can occur during catch-up processing, online forward recovery (OFR), or auto-archive.

If only a single copy of the SLDS exists (or an error occurs on both copies), the current operation is terminated, and the record of the SLDS is deleted from the RECON data set. Any missing log data is obtained from the isolated log sender.

When dual SLDSs exist, the second copy is opened and processing proceeds (assuming no error occurs on the second copy). As long as successive blocks can be read from one of the data sets, the operation proceeds.

In either case, an error message is issued. The operator must restart the OFR processing, which was interrupted by the read error, if necessary.

**Related reading:** For more information on restarting the OFR processing, see *IMS Version 11 Operations and Automation*.

- Write error on an archive data set (SLDS or RLDS)

Write errors during archive cause termination of the current archive operation. Eventually a new archive operation is initiated, and the log router retries the failed archive.

#### **DL/I database tracking**

DL/I database tracking can tolerate some unexpected return codes from other RSR components, such as from DBRC, as it tries to obtain database authorization. When it receives an unexpected error, database tracking issues a message that identifies the resource (such as the database name), and the error or unexpected return code.

Rather than bring down the entire tracking IMS, database tracking then initiates "STOP" processing for the specific database, the smallest appropriate unit of resource associated with the error. "STOP" processing includes the termination of tracking for the database, closing the database, and unauthorizing and deallocating the database.

After an error in a specific database, tracking continues for other databases that are to be tracked.

#### **Fast Path database tracking**

Fast Path database tracking can tolerate some unexpected return codes from other RSR components, such as from DBRC, as it tries to obtain area authorization. When it receives an unexpected error, Fast Path database tracking issues a message that identifies the resource (such as the area name), and the error or unexpected return code.

Rather than bring down the entire tracking IMS, Fast Path database tracking then initiates "STOP" processing for the specific area, the smallest appropriate unit of resource associated with the error. "STOP" processing includes the termination of tracking for the area, elimination of associated log records from the data space, and unauthorizing and deallocating the area.

After an error in a specific area, tracking continues for other areas that are to be tracked.

The Fast Path database tracking IMS handles I/O errors as follows:

- **Read Errors**

Read activity is done in READ ANY mode. If the record can be read from at least one ADS, the Fast Path database tracking IMS is not aware of the read error. If the record cannot be read from any ADS, the Fast Path database tracking IMS stops tracking the area.

If you have an error in one ADS, copy the other one (which has no error) immediately, then discard the one with the error.

- **Write Errors**

An error queue element (EQE) is created for the control interval in error, just as is done in a non-RSR environment.

When additional ADSs need to be added to the existing area, do the following:

1. Stop the tracked area, using the `/DBR AREA areaname` command.
2. Create additional ADSs, using the Access Method Services `REPRO` command.
3. Register new ADSs with DBRC, using the `INIT.ADS` and `CHANGE.ADS AVAIL` commands.
4. Start online forward recovery for the area, using the `/START AREA areaname` command.

**Note:** The ADS CREATE utility can not be used at the remote site because the remote site has no dependent regions.

#### **Online forward recovery**

When online forward recovery (OFR) for shadowed databases and areas is stopped or fails, the point at which OFR stopped on the SLDS is recorded for each database and area undergoing online forward recovery. This point becomes the “restart position” for OFR when the database or area starts again.

#### **Online change**

I/O errors that occur when attempting to read or write to a tracking MODSTAT data set do not cause the tracking IMS to fail. Information about the error is saved so that the read or write can be retried. The read or write is retried when the next checkpoint or online change record (X'70') is received.

Tracking MODSTAT error messages are issued when the error first occurs and when each online change record is received. Error information is saved across tracking IMS restarts. All failed reads and writes are retried during tracking IMS restarts, and appropriate messages are issued.

---

## **Establishing IMS security for RSR**

The considerations for establishing IMS security relative to RSR are Transport Manager Subsystem (TMS) and IMS terminal security.

#### **Transport Manager Subsystem (TMS)**

The TMS provides basic security by allowing you to define an ACB password to secure the access to the TMS APPLID. You need to protect against unauthorized access to libraries with VTAM definitions by carefully using RACF or an equivalent security product.

The TMS only communicates with a logical unit (LU) if that LUname has been specified on a DEFINE SYSTEM command. Combined with network security, this prevents TMS from communicating with unauthorized programs or terminals.

All users of the TMS must be authorized programs; that is, they must be APF-authorized or must run in system supervisor state.

#### **IMS terminal security**

It is important to keep terminal-related passwords synchronized between the active and remote sites.



---

## Chapter 46. Recovery in an IMSplex

Recovery in an IMSplex is performed differently for the member IMS systems and Common Service Layer (CSL).

The following topics provide an overview of each type of recovery:

### Recovery of IMS systems in an IMSplex

The recovery procedures for a failed IMS system within an IMSplex are the same as the procedures for recovering a stand-alone IMS system. For example, the procedures for an IMS system that terminates because of a z/OS failure or a hardware power failure are the same regardless of whether the IMS system is a member of an IMSplex: you must IPL z/OS and then restart IMS using an /ERESTART command.

### Recovery of CSL in an IMSplex

If an entire CSL fails, each CSL manager must be restarted. If OM fails and you only have one OM defined for an IMSplex, you lose your type-2 command interface capability and cannot enter commands through a SPOC. You can still enter type-1 commands through supported sources. OM can be restarted by using the z/OS Automatic Restart Manager (ARM) as defined on the OM execution parameter, specifying a started procedure, or submitting JCL.

Similarly, if RM or SCI fails, each must be restarted, either by the z/OS Automatic Restart Manager, specifying a started procedure, or submitting JCL. If no RMs are available, the IMSplex continues functioning, but all RM functions cease until an RM is restarted. If every RM, resource structure, coupling facility, or CQS fails or is inaccessible within an IMSplex, global resource management does not function until they are all restarted. If every SCI fails, the CSL cannot operate because all communication is down. SCI failure also affects automatic RECON loss notification. SCI must be restarted to bring the CSL back in operation.





---

## Part 5. Using IMS reports

These topics describe IMS reports and how to use them.



---

## Chapter 47. DB Monitor reports

These topics describe the DB-Monitor reports and how to use them.

**Related concepts:**

 Monitoring VSAM buffers (Database Administration)

---

### VSAM-Buffer-Pool report

The VSAM-Buffer-Pool report gives you processing information about a specific VSAM subpool. One report is produced for each subpool you specify.

For subpools within a VSAM local shared resource pool, reports are produced in ascending order based on subpool buffer size. If both index and data subpools exist in a given shared resource pool, index subpool reports follow data subpool reports. Reports for subpools in different shared resource pools are always produced in the order the shared resource pools are specified. You specify shared resource pools and subpools in control statements processed when IMS is initialized. In a batch system, the control statements are put in the DFSVSAMP data set. In an online system, they are put in the IMS.PROCLIB data set with the DFSVSMnn member name.

The VSAM-Buffer-Pool report has no meaning for HSAM or SHSAM databases, because neither of these databases can use VSAM as the access method.

### Using the VSAM-Buffer-Pool report

The primary usefulness of the VSAM-Buffer-Pool report is to calculate how many I/O operations were required to read to or write from the subpool.

You might want to increase buffer pool size to see if you can decrease the number of I/O operations. You might also want to turn on background write. Or, if the number of I/O operations is increasing over time, you might need to reorganize your database.

To calculate the number of I/O operations required to read to or write from the buffer pool, use the following formula:

Total I/O equals the sum of:

1. The number of times a CI had to be read into the buffer from the database (NUMBER of CONTR INT READ FROM EXTERNAL STORAGE field in the report).
2. The number of times a buffer had to be written to the database (NUMBER OF VSAM WRITES INITIATED BY IMS field in the report).
3. The number of times a buffer had to be written to the database so a new CI could be read into the buffer (NUMBER OF VSAM WRITES TO MAKE SPACE IN THE POOL field in the report).

### Fields in the VSAM-Buffer-Pool report

The following example is a sample of a VSAM-Buffer-Pool report.

V S A M   B U F F E R   P O O L

FIX INDEX/BLOCK/DATA  
SHARED RESOURCE POOL ID

N/Y/N  
VPL/1

|                                                        |                           |           |            |
|--------------------------------------------------------|---------------------------|-----------|------------|
|                                                        | SHARED RESOURCE POOL TYPE |           | D          |
|                                                        | SUBPOOL ID                |           | 1          |
|                                                        | SUBPOOL BUFFER SIZE       |           | 2048       |
|                                                        | NUMBER HIPERSPACE BUFFERS |           | 0          |
|                                                        | TOTAL BUFFERS IN SUBPOOL  |           | 4          |
|                                                        | 17:08:15                  | 17:10:16  |            |
|                                                        | START TRACE               | END TRACE | DIFFERENCE |
| NUMBER OF RETRIEVE BY RBA CALLS RECEIVED BY BUF HNDLR  | 0                         | 0         | 0          |
| NUMBER OF RETRIEVE BY KEY CALLS                        | 0                         | 0         | 0          |
| NUMBER OF LOGICAL RECORDS INSERTED INTO ESDS           | 0                         | 0         | 0          |
| NUMBER OF LOGICAL RECORDS INSERTED INTO KSDS           | 0                         | 0         | 0          |
| NUMBER OF LOGICAL RECORDS ALTERED IN THIS SUBPOOL      | 0                         | 0         | 0          |
| NUMBER OF TIMES BACKGROUND WRITE FUNCTION INVOKED      | 0                         | 0         | 0          |
| NUMBER OF SYNCHRONIZATION CALLS RECEIVED               | 18566                     | 27923     | 9357       |
| NUMBER OF WRITE ERROR BUFFERS CURRENTLY IN THE SUBPOOL | 0                         | 0         | 0          |
| LARGEST NUMBER OF WRITE ERRORS IN THE SUBPOOL          | 0                         | 0         | 0          |
| NUMBER OF VSAM GET CALLS ISSUED                        | 0                         | 0         | 0          |
| NUMBER OF VSAM SCHBFR CALLS ISSUED                     | 0                         | 0         | 0          |
| NUMBER OF TIMES CONT INT REQUESTED ALREADY IN POOL     | 229956                    | 349375    | 119419     |
| NUMBER OF CONTR INT READ FROM EXTERNAL STORAGE         | 1078                      | 1229      | 151        |
| NUMBER OF VSAM WRITES INITIATED BY IMS                 | 33                        | 64        | 31         |
| NUMBER OF VSAM WRITES TO MAKE SPACE IN THE POOL        | 0                         | 0         | 0          |
| NUMBER OF VSAM READS FROM HIPERSPACE BUFFERS           | 0                         | 0         | 0          |
| NUMBER OF VSAM WRITES FROM HIPERSPACE BUFFERS          | 0                         | 0         | 0          |
| NUMBER OF FAILED VSAM READS FROM HIPERSPACE BUFFERS    | 0                         | 0         | 0          |
| NUMBER OF FAILED VSAM WRITES FROM HIPERSPACE BUFFERS   | 0                         | 0         | 0          |

The meaning of the various fields in the report is as follows:

#### **FIX INDEX/BLOCK/DATA**

This field indicates the fix options for the index buffers/data buffer prefix/data buffers for this subpool.

#### **SHARED RESOURCE POOL ID**

This field is a four character, VSAM local-shared-resource-pool ID provided at subpool definition time. This ID is specified in the DFSVSAMP or IMS.PROCLIB data set control statements.

#### **SHARED RESOURCE POOL TYPE**

This field indicates whether this subpool contains index (I) or data (D) buffers.

#### **SUBPOOL ID**

This field describes the subpool ID. A unique subpool ID is assigned for each different database buffer size in each VSAM local shared resource pool you specify in the DFSVSAMP or IMS.PROCLIB data set control statements.

#### **SUBPOOL BUFFER SIZE**

This field describes the size, in bytes, of buffers in this subpool. You can specify buffers of different sizes. A subpool contains all buffers of the same size. All subpools grouped together make up the buffer pool. Buffer size is specified by you in the control statements for the DFSVSAMP or IMS.PROCLIB data sets. For example, suppose you specified:

```
//DFSVSAMP DD input for VSAM and OSAM buffers and options
:
:
POOLID=BB
VSRBF=4096,4,D,HSR,10
VSRBF=8192,4,D,HSO,20
POOLID=FF
VSRBF=4096,4,I
VSRBF=8192,4,D
:
/*
```

This gives you four subpools and the four VSAM-Buffer-Pool reports displayed in the following examples. The headings in each report display as shown in the examples.

```

SHARED RESOURCE POOL ID:    BB
SHARED RESOURCE POOL TYPE:  D
SUBPOOL ID:                 1
SUBPOOL BUFFER SIZE:        4096
HIPERSPACE BUFFERS:         10
TOTAL BUFFERS IN SUBPOOL:   4

SHARED RESOURCE POOL ID:    BB
SHARED RESOURCE POOL TYPE:  D
SUBPOOL ID:                 2
SUBPOOL BUFFER SIZE:        8192
HIPERSPACE BUFFERS:         20
TOTAL BUFFERS IN SUBPOOL:   4

SHARED RESOURCE POOL ID:    FF
SHARED RESOURCE POOL TYPE:  D
SUBPOOL ID:                 1
SUBPOOL BUFFER SIZE:        8192
HIPERSPACE BUFFERS:         0
TOTAL BUFFERS IN SUBPOOL:   4

SHARED RESOURCE POOL ID:    FF
SHARED RESOURCE POOL TYPE:  I
SUBPOOL ID:                 2
SUBPOOL BUFFER SIZE:        4096
HIPERSPACE BUFFERS:         0
TOTAL BUFFERS IN SUBPOOL:   4

```

#### NUMBER OF HIPERSPACE BUFFERS

This field indicates the number of hiperspace buffers specified at subpool definition time.

#### TOTAL BUFFERS IN SUBPOOL

This field describes how many buffers are in the specified subpool.

#### START TRACE, END TRACE, and DIFFERENCE

The start trace and end trace fields tell you the time the DB Monitor program was *last* started and stopped. The time is generated by the time-of-day clock. Clock times are read as follows:

Clock time = hh.mm.ss

where:

**hh**      Hours 0 through 23

**mm**      Minutes

**ss**      Seconds

If the DB Monitor program was on during an entire batch run, the start and end trace times are when the batch run started and stopped. If the DB Monitor program was turned on and off more than once in the same batch run, the start and end trace times are when the monitor was *last* started and stopped.

The numbers under the start and end trace fields are cumulative numbers for the entire batch run. If the monitor was only turned on and off once, the start trace number is zero. If the monitor was turned on and off more than once, the start trace numbers are the cumulative numbers when the monitor was last turned on; the stop trace numbers are the cumulative numbers when the monitor was last turned off.

The difference column describes the difference between the cumulative numbers in the start and end trace fields. If the monitor was only turned on and off once, the difference column contains the same numbers as the end trace column.

#### **NUMBER OF RETRIEVE BY RBA CALLS**

This field describes how many retrieve by relative byte address (RBA) calls were issued for this subpool. Retrieve by RBA calls are calls issued internally by DL/I. One retrieve by RBA call is issued for each direct-address pointer that must be followed in searching for a segment. For example, a GN call for a dependent segment in an HDAM or PHDAM database uses a series of RBA calls to search for the dependent segment, one call for each direct-address pointer it follows.

If you want to know the exact sequence of a search when a retrieve by RBA call is used, you can record the sequence by turning on the buffer handler trace and using a SNAP call to see the trace records. You can turn on the buffer handler trace using the DL/I= operand on the OPTIONS statement for the DFSVSAMP or DFSVSMnn data set. The SNAP call can be issued from the application program or by using the DFSDDLTO test program.

One call from an application program can generate more than one retrieve by RBA call. The retrieve by RBA call might or might not require an I/O operation. Because the number in this field does not reflect the number of I/O operations to access a segment, do not use it to judge VSAM performance.

#### **NUMBER OF RETRIEVE BY KEY CALLS**

This field describes how many retrieve by key calls were issued for this subpool. Retrieve by key calls are calls issued internally by DL/I. The calls are issued to search a KSDS using a key as a qualification (where key is equal to or greater than X). For example, a GU call for a root segment in a HIDAM database causes DL/I to issue a retrieve by key call to access the index segment pointing to the requested root segment.

One call from an application program can generate more than one retrieve by key calls. The retrieve by key calls might or might not require an I/O operation. Because the number in this field does not reflect the number of I/O operations to access a segment, do not use it to judge VSAM performance.

#### **NUMBER OF LOGICAL RECORDS INSERTED INTO ESDS**

This field describes how many of the logical records in your ESDS that were previously empty now contain segments. When a dependent segment is inserted into an ESDS in a HISAM or HIDAM database, the segment might not fit into a logical record that already contains other segments. In this case, the segment is put into a new ESDS logical record. When a dependent segment is inserted into a logical record in an ESDS in a HISAM database, other segments in the same logical might need to be shifted into a new ESDS logical record to make room for the segment being inserted.

Look at this field from one report to the next. It helps you determine when you are running out of logical records in the primary space you have allocated. It is best to avoid using logical records from secondary space because this space is probably not close to the primary space. The distance between the two areas of space might cause extra seek time and therefore poor performance. In general, it is best to reorganize your database before you need to use secondary space.

#### **NUMBER OF LOGICAL RECORDS INSERTED INTO KSDS**

This field describes how many of the logical records in your KSDS that were previously empty now contain segments. HISAM databases use a new logical record when a root segment is inserted. HIDAM index databases use a new logical record for the index segment created when a root segment is inserted. Secondary index databases use a new logical record when a new pointer segment is inserted.

Look at this field from one report to the next. It helps you determine when you are running out of logical records in the primary space you have allocated. It is best to avoid using logical records from secondary space because this space is probably not close to the primary space. The distance between the two areas of space might cause extra seek time and therefore poor performance. In general, it is best to reorganize your database before you need to use secondary space.

#### **NUMBER OF LOGICAL RECORDS ALTERED IN THIS SUBPOOL**

This field describes how many logical records, while in the buffer pool, were marked as altered. When a segment is inserted or replaced in a logical record, the logical record in the buffer is marked as altered until it is written back to the database.

#### **NUMBER OF TIMES BACKGROUND WRITE FUNCTION INVOKED**

If you have specified use of the background write function, this field tells how many times the function was used. The background write function, at intervals, writes buffers containing modified data back to the database. It does this so buffers are available for use when an application program needs them. Without background write, if an application program wants to read data into a buffer that already contains modified data, the application program has to wait while the contents of the buffer are written back to the database. The number of times background write is invoked is the same on each subpool report produced, during a given execution of the monitor, for a given local shared resource pool. Once invoked, the background write function writes buffers from all subpools within a local shared resource pool.

Background write is specified in the BGWRT= operand of the OPTIONS statement for the DFSVSAMP or DFSVSMnn data set.

#### **NUMBER OF SYNCHRONIZATION CALLS RECEIVED**

This field describes how many requests were made to write all altered buffers in the local shared resource pool while the monitor was on. CHKP and STAT calls are typical requestors of this.

#### **NUMBER OF PERM WRT ERROR BUFFS NOW IN THE SUBPOOL**

This field describes how many buffers are currently “frozen” in storage because a permanent I/O error occurred when writing them to the database. When a VSAM write operation results in a permanent I/O error, the affected buffers are frozen in storage until the data set is closed or, in the online system, until the system is shut down. Once the data set is closed or the system shut down, the buffers are written to the log and the number in this field returns to zero.

Ensure that the operator performed database recovery of the affected data set before you ever receive this report.

#### **LARGEST NUMB OF PERM ERR BUFFS EVER IN THE SUBPOOL**

This field describes how many buffers *were* frozen in storage when the condition described in the previous field occurred.

#### **NUMBER OF VSAM GET CALLS ISSUED**

This field describes how many times VSAM GET calls were issued. VSAM GET calls are calls issued internally by DL/I. The GET call might be satisfied by data in the buffer pool or it might require that data be read into the buffer pool. Because the number in this field does not reflect the number of I/O operations required to access a segment, do not use it to judge VSAM performance.



**NUMBER OF VSAM SCHBFR CALLS ISSUED**

This field describes how many times the HD space management routine issued calls to search for space in which to insert segments.

If, from one monitor report to the next, the number in this field is increasing, it means that space for storing new segments is not available in the most desirable location. Eventually, this means you must reorganize your database to improve performance. In reorganizing, pay special attention to the operands affecting database space (the bytes operand in the RMNAME= keyword in the DBD statement and the fbff and fspf operands in the FRSPC= keyword in the DATA SET statement).

**NUMBER OF TIMES CONTR INT REQUESTED ALREADY IN POOL**

This field describes how many times a logical record was found in a CI that was already in the buffers. When this occurs, no I/O operations are required to access the desired segments.

If you are trying to improve performance, increase the number of buffers you have allocated. If you increase the number of buffers, you can monitor this field to see if the number in it increases, which indicates improved performance.

**NUMBER OF CONTR INT READ FROM EXTERNAL STORAGE**

This field describes how many times a logical record was not found in a CI that was already in the buffers. When this occurs, an I/O operation is required to read the CI containing the logical record into the buffer pool. Because performance is always better when fewer I/O operations are performed, you might want to increase the number of buffers you have specified to see how that affects the number in this field. Specifying more buffers keeps more CIs (and therefore logical records) in the buffer pool. There is a break-even point in this process, however, where too many buffers are specified, and it takes longer to search and maintain the buffers than it takes to read a CI into the buffer.

The number of buffers is specified in the control statements for the DFSVSAMP or DFSVSMnn data sets.

**NUMBER OF VSAM WRITES INITIATED BY IMS**

This field describes the number of times DL/I issued a write request to write data to the database. Write operations are issued when:

- A data set is closed. Database buffers containing data that has been altered by the data set being closed are written to the database.
- Abnormal termination occurs during application program processing. Database buffers containing data that has been altered are written to the database.
- The background write function is invoked. Selected database buffers containing data that has been altered are written to the database.
- A checkpoint call is issued. All altered database buffers are written to the database.

**NUMBER OF VSAM WRITES TO MAKE SPACE IN THE POOL**

This field describes how many times a CI had to be read into a buffer containing a logical record with altered data. When this happens, the buffer (because it contains altered data) has to be written back to the database before the new CI can be read into it. This means the application program has to wait while the write operation takes place.

For best performance, ensure that the number in this field is close to zero. This can generally be achieved by turning on the background write function during batch processing and adjusting the number of buffers allocated.

#### NUMBER OF VSAM READS FROM HIPERSPACE BUFFERS

This field describes the total number of successful VSAM reads (MOVEPAGE and NON-MOVEPAGE) from the hiperspace buffers.

#### NUMBER OF VSAM WRITES FROM HIPERSPACE BUFFERS

This field describes the total number of successful VSAM writes (MOVEPAGE and NON-MOVEPAGE) to the hiperspace buffers.

#### NUMBER OF FAILED VSAM READS FROM HIPERSPACE BUFFERS

This field describes the number of times that a VSAM read request from hiperspace failed, resulting in a read from DASD.

#### NUMBER OF FAILED VSAM WRITES FROM HIPERSPACE BUFFERS

This field describes the number of times that a VSAM WRITE request to hiperspace failed, resulting in a write to DASD.

## VSAM-Statistics report

VSAM-Statistics reports are based on: a specific application program PCB, a data set the application program is using, and the type of DL/I call the application program issued.

The VSAM-Statistics report has no meaning for HSAM or SHSAM databases because neither of these databases can use VSAM as the access method.

## Using the VSAM-Statistics report

From a VSAM-Statistics report, you can determine which calls in an application program require a great many I/O operations. After you know this, you can improve performance by tuning either the database or the application program to reduce I/O operations.

The following fields in the report tell actual I/O activity and are therefore the most important ones to monitor:

- READS
- USR WTS
- NUR WTS

If you can reduce the averages in these fields, performance can be improved.

In tuning to reduce I/O operations, pay most attention to calls issued a large number of times. It is more profitable to save 1 second on a call executed 2000 times than to save 5 seconds on a call executed ten times (2000 versus 50 seconds). The DL/I-Call-Summary report (discussed under “DL/I-Call-Summary report” on page 724) describes how many times each DL/I application program call is issued.

## Fields in the VSAM-Statistics report

The following example is a sample of a VSAM-Statistics report.

```

IMS MONITOR  ****VSAM STATISTICS****      TRACE START 1989 076 12:42:54      TRACE STOP 1989 076 12:43:07  PAGE 0012
PCBNAME  DDNAME  DL/I  VSAM  RET  RET  ISRT  ISRT  BFR  BKG  SYN  GETS  SCHBFR  FOUND  READS  USR  NUR
FUNC  IWAITS  RBA  KEY  ESDS  KSDS  ALT  WTS  PTS  WTS  WTS  WTS  WTS  WTS  WTS  WTS  WTS
DLVNTZ02 DBHVSAM2 STAT 1 1.00 0.00 0.00 0.00 1.00 0.00 1.00 1.00 0.00 1.00 0.00 1.00 0.00
DD TOTAL

```

|             |          |      |    |        |       |       |        |       |      |      |        |      |        |      |      |      |
|-------------|----------|------|----|--------|-------|-------|--------|-------|------|------|--------|------|--------|------|------|------|
|             |          |      | 1  | 1.00   | 0.00  | 0.00  | 0.00   | 1.00  | 0.00 | 1.00 | 1.00   | 0.00 | 1.00   | 0.00 | 1.00 | 0.00 |
|             | HIDAM    | STAT | 4  | 307.25 | 36.50 | 0.50  | 0.00   | 35.00 | 0.00 | 0.25 | 120.25 | 0.00 | 157.00 | 0.25 | 0.75 | 0.00 |
|             | DD TOTAL |      |    |        |       |       |        |       |      |      |        |      |        |      |      |      |
|             | XDLBT04I | GU   | 4  | 307.25 | 36.50 | 0.50  | 0.00   | 35.00 | 0.00 | 0.25 | 120.25 | 0.00 | 157.00 | 0.25 | 0.75 | 0.00 |
|             | DD TOTAL |      | 2  | 15.00  | 4.50  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 17.50  | 0.00 | 21.50  | 1.00 | 0.00 | 0.00 |
| PCB TOTAL   |          |      | 2  | 15.00  | 4.50  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 17.50  | 0.00 | 21.50  | 1.00 | 0.00 | 0.00 |
|             |          |      | 7  | 180.00 | 22.14 | 0.28  | 0.00   | 20.14 | 0.00 | 0.28 | 73.85  | 0.00 | 96.00  | 0.42 | 0.57 | 0.00 |
| DLVNTZX2    | HIDAM    | GN   | 1  | 15.00  | 3.00  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 18.00  | 0.00 | 20.00  | 1.00 | 0.00 | 0.00 |
|             | GU       |      | 1  | 1.00   | 0.00  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 1.00   | 0.00 | 0.00   | 1.00 | 0.00 | 0.00 |
|             | DD TOTAL |      |    |        |       |       |        |       |      |      |        |      |        |      |      |      |
|             | DBHVSAM1 | GU   | 2  | 8.00   | 1.50  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 9.50   | 0.00 | 10.00  | 1.00 | 0.00 | 0.00 |
|             | DD TOTAL |      | 8  | 0.00   | 0.12  | 12.75 | 520.00 | 7.50  | 0.00 | 0.00 | 0.12   | 0.00 | 0.00   | 0.87 | 0.00 | 0.00 |
|             |          |      | 8  | 0.00   | 0.12  | 12.75 | 520.00 | 7.50  | 0.00 | 0.00 | 0.12   | 0.00 | 0.00   | 0.87 | 0.00 | 0.00 |
|             | XDLBT04I | GU   | 6  | 0.00   | 0.00  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 0.00   | 0.00 | 0.00   | 1.00 | 0.00 | 0.00 |
|             | DD TOTAL |      |    |        |       |       |        |       |      |      |        |      |        |      |      |      |
|             | DBHVSAM2 | GU   | 6  | 0.00   | 0.00  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 0.00   | 0.00 | 0.00   | 1.00 | 0.00 | 0.00 |
|             | DD TOTAL |      | 3  | 0.66   | 0.00  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 0.66   | 0.00 | 0.00   | 1.00 | 0.00 | 0.00 |
|             |          |      | 3  | 0.66   | 0.00  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 0.66   | 0.00 | 0.00   | 1.00 | 0.00 | 0.00 |
| PCB TOTAL   |          |      | 3  | 0.66   | 0.00  | 0.00  | 0.00   | 0.00  | 0.00 | 0.00 | 0.66   | 0.00 | 0.00   | 1.00 | 0.00 | 0.00 |
|             |          |      | 19 | 0.94   | 0.21  | 5.36  | 482.10 | 3.15  | 0.00 | 0.00 | 1.15   | 0.00 | 1.05   | 0.94 | 0.00 | 0.00 |
| BATCH TOTAL |          |      | 26 | 49.15  | 6.11  | 3.84  | 83.07  | 3.11  | 0.00 | 0.07 | 20.73  | 0.00 | 26.61  | 0.80 | 0.15 | 0.00 |

The meaning of the various fields in the report is as follows:

#### TRACE START and TRACE STOP

The trace start and trace stop fields tell you the time when the DB Monitor program was last started and stopped. The time is generated by the time-of-day clock. Clock times are read as follows:

Clock time = hh.mm.ss

where:

**hh** Hours 0 through 23

**mm** Minutes

**ss** Seconds

If the DB Monitor was on during an entire batch run, the trace start and trace stop time is when the batch run started and stopped. If the DB Monitor was turned on and off more than once in the same batch run, the trace start and trace stop time is when the monitor was last started and stopped.

#### PCBNAME

This field describes the name of the PCB the report is providing information about. Remember that each application program has one or more PCBs. This field can only be used to identify which application program the report is providing information about if PCB names are unique to application programs.

#### DDNAME

This field describes the name of the data set the application program is using. Within one report, an application program (PCB) can access more than one data set. The statistics compiled are listed separately for each data set.

#### DL/I FUNC

This field describes the type of DL/I calls the application program issued.

#### VSAM IWAITS

This field describes, by type of DL/I call against a specific data set, the number of times IMS had to wait before processing could proceed. When IMS has to wait, it is almost always waiting for an I/O operation to take place, that is, data is being either read from the database to the buffer or written from the buffer back to the database.

The numbers in each column under all remaining fields in the report are averages. They tell the average number of times an activity occurred rather than the specific number of times. Averages are for waits. These numbers are truncated. For example, a value of 0.019 is printed as 0.01.

#### **RET RBA**

This field describes how many retrieve by relative byte address (RBA) calls were issued for this subpool. Retrieve by RBA calls are calls issued internally by DL/I. One retrieve by RBA call is issued for each direct-address pointer that must be followed in searching for a segment. For example, a GN call for a dependent segment in an HDAM or PHIDAM database uses a series of RBA calls to search for the dependent segment, one call for each direct-address pointer it follows.

One call from an application program can generate more than one retrieve by RBA call. The retrieve by RBA call might or might not require an I/O operation. Because the number in this field does not reflect the number of I/O operations to access a segment, do not use it to judge VSAM performance.

#### **RET KEY**

This field describes how many retrieve by key calls were issued for this subpool. Retrieve by key calls are calls issued internally by DL/I. The calls are issued to search a KSDS using a key as a qualification (where key is equal to or greater than X). For example, a GU call for a root segment in a HIDAM or PHIDAM database causes DL/I to issue a retrieve by key call to access the index segment pointing to the requested root segment.

One call from an application program can generate more than one retrieve by key calls. The retrieve by key calls might or might not require an I/O operation. Because the number in this field does not reflect the number of I/O operations to access a segment, do not use it to judge VSAM performance.

#### **ISRT ESDS**

This field describes how many of the logical records in your ESDS that were previously empty now contain segments. When a dependent segment is inserted into an ESDS in a HISAM, HIDAM or PHIDAM database, the segment might not fit into a logical record that already contains other segments. In this case, the segment is put into a new ESDS logical record. When a dependent segment is inserted into a logical record in an ESDS in a HISAM database, other segments in the same logical record might need to be shifted into a new ESDS logical record to make room for the segment being inserted.

Look at this field from one report to the next. It helps you determine when you are running out of logical records in the primary space you have allocated. It is best to avoid using logical records from secondary space because this space is probably not close to the primary space.

#### **ISRT KSDS**

This field describes how many of the logical records in your KSDS that were previously empty now contain segments. HISAM databases use a new logical record when a root segment is inserted. HIDAM and PHIDAM index databases use a new logical record for the index segment created when a root segment is inserted.

Look at this field from one report to the next. It helps you determine when you are running out of logical records in the primary space you have allocated. It is best to avoid using logical records from secondary space because this space is probably not close to the primary space. The distance between the two

areas of space might cause extra seek time and therefore poor performance. In general, it is best to reorganize your database before you need to use secondary space.

#### **BFR ALT**

This field describes how many logical records, while in the buffer pool, were marked as altered. When a segment is inserted or replaced in a logical record, the logical record in the buffer is marked as altered until it is written back to the database.

#### **BKG WTS**

If you have specified use of the background write function, this field tells how many times the function was used. Background write, at intervals, writes buffers containing modified data back to the database. It does this so buffers are available for use when an application program needs them. Without background write, if an application program wants to read data into a buffer that already contains modified data, the application program has to wait while the contents of the buffer are written back to the database. The number of times background write was invoked is the same on each subpool report produced during a given execution of the monitor. This is because, once involved, background write writes buffers from all subpools.

Background write is specified in the BGWRT= operand of the OPTIONS statement for the DFSVSAMP or DFSVSMnn data set.

#### **SYN PTS**

This field describes how many times checkpoint calls were issued in DL/I programs while the monitor was on.

#### **GETS**

This field describes how many times VSAM GET calls were issued. VSAM GET calls are calls issued internally by DL/I. The GET call might be satisfied by data in the buffer pool or it might require that data be read into the buffer pool. Because the number in this field does not reflect the number of I/O operations required to access a segment, do not use it to judge VSAM performance.

#### **SCHBFR**

This field describes how many times the HD space management routine issued calls to search for space in which to insert segments.

If, from one monitor report to the next, the number in this field is increasing, it means that space for storing new segments is not available in the most desirable location. Eventually, you must reorganize your database to improve performance. In reorganizing, pay special attention to the operands affecting database space (the BYTES operand in the RMNAME= keyword in the DBD statement and the fbff and fspf operands in the FRSPC= keyword in the DATA SET statement).

#### **FOUND**

This field describes how many times a logical record was found in a CI that was already in the buffers. When this occurs, no I/O operations are required to access the desired segments.

If you are trying to improve performance, increase the number of buffers you have allocated. If you increase the number of buffers, you can monitor this field to see if the number in it increases, which indicates improved performance.

#### **READS**

This field describes how many times a logical record was not found in a CI

that was already in the buffers. When this occurs, an I/O operation is required to read the CI containing the logical record into the buffer pool. Because performance is always better when fewer I/O operations are performed, you might want to increase the number of buffers you have specified to see how that affects the number in this field. Specifying more buffers keeps more CIs (and therefore logical records) in the buffer pool. There is a break-even point in this process, however, where too many buffers are specified, and it takes longer to search and maintain the buffers than it takes to read a CI into the buffer.

The number of buffers is specified in the control statements for the DFSVSAMP or DFSVSMnn data sets.

#### **USR WTS**

This field, which indicates user writes, describes the number of times DL/I issued a write request to write data to the database. Write operations are issued when:

- A data set is closed. Database buffers containing data that has been altered by the data set being closed are written to the database.
- Abnormal termination occurs during application program processing. Database buffers containing data that has been altered are written to the database.
- The background write function is invoked. Selected database buffers containing data that has been altered are written to the database.
- A checkpoint call is issued. All altered database buffers are written to the database.

#### **NUR WTS**

This field, which indicates nonuser writes, tells how many times a CI had to be read into a buffer containing a logical record with altered data. When this happens, the buffer (because it contains altered data) has to be written back to the database before the new CI can be read into it. This means the application program has to wait while the write operation takes place.

For best performance, ensure that the number in this field is close to zero. This can generally be achieved by turning on the background write function during batch processing and adjusting the number of buffers allocated.

#### **SCRS**

This field describes the total number of successful VSAM reads (MOVEPAGE and NON-MOVEPAGE) from hiperspace buffers.

#### **SCWS**

This field describes the total number of successful VSAM writes (MOVEPAGE and NON-MOVEPAGE) to hiperspace buffers.

#### **SCRF**

This field describes the number of times that a VSAM read request from hiperspace failed, resulting in a read from DASD.

#### **SCWF**

This field describes the number of times that a VSAM write request to hiperspace failed, resulting in a write to DASD.

#### **DD TOTAL**

This field describes, for a given data set, the overall average number of times an activity occurred (except for the VSAM CALLS field, which tells total number of times).

**PCB TOTAL**

This field describes, for a given PCB, the overall average number of times an activity occurred (except for the VSAM CALLS field, which tells total number of times).

**BATCH TOTAL**

This field describes, for the time the monitor was running, the overall average number of times an activity occurred (except for the VSAM CALLS field, which tells total number of times).

---

## Database-Buffer-Pool report

The Database-Buffer-Pool report gives you information about OSAM subpools during processing. One report is produced for all subpools in the buffer pool.

This report and the VSAM-Buffer-Pool report differ in that the VSAM report was produced for each subpool in the buffer pool.

The Database-Buffer-Pool report has no meaning for HSAM, SHSAM, SHISAM, or GSAM databases because none of these databases can use OSAM as the access method.

## Using the Database-Buffer-Pool report

The primary usefulness of the Database-Buffer-Pool report is to calculate how many I/O operations were required to read to or write from the OSAM buffer pool.

You might want to increase buffer pool size to see if you can decrease the number of I/O operations. Or, if the number is increasing over time, you might need to reorganize your database.

To calculate the number of I/O operations required to read to or write from the buffer pool, use the following formula:

Total I/O equals the sum of:

1. Blocks read from the database  
Number of blocks read for OSAM specific requests (Number of Read Requests Issued)
2. Blocks written to the database  
Number of blocks written because of buffer steal processing and purge processing (Number of Blocks Written)

## Fields in the Database-Buffer-Pool report

The following figure is an example of a Database Buffer Pool report.

| DATA BASE BUFFER POOL                                  |          |                          |        |
|--------------------------------------------------------|----------|--------------------------|--------|
|                                                        |          | FIX PREFIX/BUFFERS       | Y/Y    |
|                                                        |          | SUBPOOL ID               | 004K   |
|                                                        |          | SUBPOOL BUFFER SIZE      | 4096   |
|                                                        |          | TOTAL BUFFERS IN SUBPOOL | 1000   |
|                                                        | 17:08:15 | 17:10:16                 |        |
| NUMBER OF LOCATE-TYPE CALLS                            | 1117674  | 1676213                  | 558539 |
| NUMBER OF REQUESTS TO CREATE NEW BLOCKS                | 0        | 0                        | 0      |
| NUMBER OF BUFFER ALTER CALLS                           | 215874   | 322936                   | 107062 |
| NUMBER OF PURGE CALLS                                  | 25077    | 37454                    | 12377  |
| NUMBER OF LOCATE-TYPE CALLS, DATA ALREADY IN OSAM POOL | 870306   | 1301187                  | 430881 |
| NUMBER OF BUFFERS SEARCHED BY ALL LOCATE-TYPE CALLS    | 1258247  | 1886843                  | 628596 |



|                                                        |        |        |        |
|--------------------------------------------------------|--------|--------|--------|
| NUMBER OF READ I/O REQUESTS                            | 238165 | 360260 | 122095 |
| NUMBER OF SINGLE BLOCK WRITES BY BUFFER STEAL ROUTINE  | 0      | 0      | 0      |
| NUMBER OF BLOCKS WRITTEN BY PURGE                      | 95057  | 142413 | 47356  |
| NUMBER OF LOCATE CALLS WAITED DUE TO BUSY ID           | 780    | 1297   | 517    |
| NUMBER OF LOCATE CALLS WAITED DUE TO BUFFER BUSY WRT   | 0      | 0      | 0      |
| NUMBER OF LOCATE CALLS WAITED DUE TO BUFFER BUSY READ  | 0      | 0      | 0      |
| NUMBER OF BUFFER STEAL/PURGE WAITED FOR OWNERSHIP RLSE | 178    | 261    | 83     |
| NUMBER OF BUFFER STEAL REQUESTS WAITED FOR BUFFERS     | 0      | 0      | 0      |
| TOTAL NUMBER OF I/O ERRORS FOR THIS SUBPOOL            | 0      | 0      | 0      |
| NUMBER OF BUFFERS LOCKED DUE TO WRITE ERRORS           | 0      | 0      | 0      |

The meaning of the various fields in the report is as follows:

#### **FIX PREFIX/BUFFERS**

This field indicates the fix options for the buffer prefix/data buffers for this subpool.

#### **SUBPOOL ID**

This field is a 4 character pool ID provided at subpool definition time.

#### **SUBPOOL BUFFER SIZE**

This field indicates the size, in bytes, of the buffers in this subpool.

#### **TOTAL BUFFERS IN SUBPOOL**

This field indicates the total number of buffers in this subpool.

On the following line are time entries that indicate the start trace and end trace times. The start trace and end trace fields tell you the time when the DB Monitor program was last started and stopped.

#### **NUMBER OF LOCATE-TYPE CALLS**

This field indicates the number of locate-type calls for this subpool.

#### **NUMBER OF REQUESTS TO CREATE NEW BLOCKS**

This field indicates the number of times a block had a segment inserted for the first time. When this happens, the block is marked as modified and must eventually be written back to the database.

#### **NUMBER OF BUFFER ALTER CALLS**

This field indicates the number of buffer alter calls for this subpool. This count includes NEW BLOCK and BYTALT calls.

#### **NUMBER OF PURGE CALLS**

This field indicates the number of purge requests for this subpool.

#### **NUMBER OF LOCATE-TYPE CALLS, DATA ALREADY IN SUBPOOL**

This field indicates the number of locate-type calls for this subpool where the data was already in an OSAM pool.

#### **NUMBER OF BUFFERS SEARCHED BY ALL LOCATE-TYPE CALLS**

This field indicates the number of buffers searched by all locate-type calls for this subpool.

#### **NUMBER OF READ I/O REQUESTS**

This field indicates the number of read I/O requests for this subpool.

#### **NUMBER OF SINGLE BLOCK WRITES BY BUFFER STEAL ROUTINE**

This field indicates the number of single block writes initiated by buffer steal routine for this subpool.

#### **NUMBER OF BLOCKS WRITTEN BY PURGE**

This field indicates the number of blocks for this subpool written by purge.

#### **TOTAL NUMBER OF I/O ERRORS FOR THIS SUBPOOL**

This field indicates the total number of I/O errors for this subpool.



#### NUMBER OF BUFFERS LOCKED DUE TO WRITE ERRORS

This field indicates how many buffers are currently “frozen” in storage because a permanent I/O error occurred when writing them to the database. When a write operation results in a permanent I/O error, the affected buffers are frozen in storage until the data set is closed or, in an online system, until the system is shut down. Once the data set is closed, or the online system shut down, the buffers are written to the log, and this number returns to 0.

#### NUMBER OF LOCATE CALLS WAITED DUE TO BUSY ID

This field indicates the number of locate calls for this subpool which waited due to busy ID.

#### NUMBER OF LOCATE CALLS WAITED DUE TO BUFFER BUSY WRT

This field indicates the number of locate calls for this subpool which waited due to buffer busy writing.

#### NUMBER OF LOCATE CALLS WAITED DUE TO BUFFER BUSY READ

This field indicates the number of locate calls for this subpool which waited due to buffer busy reading.

#### NUMBER OF BUFFER STEAL/PURGE WAITED DUE TO BUFFER BUSY READ

This field indicates the number of locate calls for this subpool which waited for ownership to be released.

#### NUMBER OF BUFFER STEAL REQUESTS WAITED FOR BUFFERS

This field indicates the number of buffer steal requests for this subpool which waited because no buffers were available to be stolen.

---

## Program-I/O report

The Program-I/O report tells how long IMS was inactive (no IMS code was being executed) because IMS was waiting for use of a resource or for completion of an event. This time is called IWAIT time and, for all practical purposes, can be considered the time IMS had to wait while an I/O operation took place.

One report is produced each time the DB Monitor is run, and the report describes IWAIT time by PCB and data set name.

All times in the Program-I/O report are in microseconds. A microsecond is one millionth of a second (in other words, 7050 microseconds equals 0.007050 seconds).

## Using the Program-I/O report

Using the Program-I/O report, you can correlate IWAIT time with a specific PCB and data set. This allows you to identify databases that are causing a relatively large number of IWAITs.

Because IWAITs are identified by PCB and data set names, you might be able to trace the large number of IWAITs back to a particular application program. If you can, give the application program special attention when tuning for performance. The DL/I-Call-Summary report, described in the following topic, helps you identify specific DL/I calls in an application program that are causing a large number of IWAITs.

## Fields in the Program-I/O report

The following example shows a sample of a Program-I/O report.

```
IMS MONITOR  ****PROGRAM I/O****          TRACE START 1989 076 12:42:54    TRACE STOP 1989 076 12:43:07    PAGE 0008
              .....IWAIT TIME.....          DISTR.
              PCB NAME      IWAITs      TOTAL      MEAN      MAXIMUM  DDNAME      MODULE      NO.
```

|             |    |        |       |        |          |     |     |
|-------------|----|--------|-------|--------|----------|-----|-----|
| DLVNTZ02    | 1  | 35853  | 35853 | 35853  | DBHVSAM2 | VBH | 127 |
|             | 4  | 257649 | 64412 | 196028 | HIDAM    | VBH | 128 |
|             | 2  | 79222  | 39611 | 62452  | XDLBT04I | VBH | 129 |
| PCB TOTAL   |    |        |       |        |          |     |     |
| DLVNTZX2    | 7  | 372724 | 53246 |        |          |     |     |
|             | 2  | 57645  | 28822 | 40686  | HIDAM    | VBH | 130 |
|             | 8  | 176622 | 22077 | 46141  | DBHVSAM1 | VBH | 131 |
|             | 6  | 105340 | 17556 | 27843  | XDLBT04I | VBH | 132 |
|             | 3  | 65296  | 21765 | 23458  | DBHVSAM2 | VBH | 133 |
| PCB TOTAL   |    |        |       |        |          |     |     |
| BATCH TOTAL | 19 | 404903 | 21310 |        |          |     |     |
|             | 26 | 777627 | 29908 |        |          |     |     |

The meaning of the various fields in the report is as follows:

#### TRACE START and TRACE STOP

The trace start and trace stop fields tell you the time when the DB Monitor program was last started and stopped. The time is generated by the time-of-day clock. Clock times are read as follows:

Clock time = hh.mm.ss

Where:

**hh** Hours 0 through 23

**mm** Minutes

**ss** Seconds

If the DB Monitor program was on during an entire batch run, the trace start and trace stop times is when the batch run started and stopped. If the DB Monitor program was turned on and off more than once in the same batch run, the trace start and trace stop times are the times at which the monitor was last started and stopped.

#### PCBNAME

This field describes the name of the PCB the report is providing information about. Remember that each application program has one or more PCBs. If PCB names are unique to application programs, this field can only be used to identify which application program the report is providing information about.

#### IWAITS

This field describes the number of times IMS was inactive.

#### IWAIT Time

This field describes the elapsed time during which IMS was inactive.

#### TOTAL

The total time IMS waited

#### MEAN

The average time IWAIT IMS waited

#### MAXIMUM

The longest single time IMS waited

#### DDNAME

This field describes the name of the data set the application program is using. Within one report, an application program (PCB) can access more than one data set. The statistics compiled are listed separately for each data set.

#### MODULE

This field describes the modules that issued the internal call (in response to a DL/I call) that caused IMS to wait.

**DBH**    Module DFSDBHR0  
**DLE**    Module DFSDDLE0  
**VBH**    Module DFSDVSM0

**DISTR.NUMBER**

The distribution number is a reference number to be used in conjunction with the Distribution-Appendix report. For a description of what it means, see “Distribution-Appendix report” on page 727.

**PCB Total**

For a given PCB this field describes:

**IWAITS**

The total number of times IMS was inactive

**TOTAL**

The total time IMS waited

**MEAN**

The average time per IWAIT

**BATCH Total**

For this execution of the DB Monitor this field describes:

**IWAITS**

The total number of times IMS was inactive

**TOTAL**

The total time IMS waited

**MEAN**

The average time per IWAIT

---

## DL/I-Call-Summary report

One DL/I-Call-Summary report is produced each time the DB Monitor is run, and the report describes times by PCB name and type of DL/I call.

The DL/I-Call-Summary report tells two things:

- How long IMS was inactive (no IMS code was being executed) because IMS was waiting for use of a resource or for completion of an event. This time is called IWAIT time and, for all practical purposes, can be considered the time IMS had to wait while an I/O operation took place.
- Of the elapsed time during which IMS was inactive, how much of it was actually IWAIT time.

All times in the DL/I-Call-Summary report are in microseconds. A microsecond is one millionth of a second (in other words, 7050 microseconds equals 0.007050 seconds).

## Using the DL/I-Call-Summary report

The primary usefulness of the DL/I-Call-Summary report is to track down DL/I calls that are causing a large number of IWAITS.

If the number in the IWAITS/CALL field is relatively high, you want to know why. Because the number in the report is related to a specific DL/I call, segment, and PCB, you can trace the IWAITS back to a specific part of an application program. In addition, by using the Distribution-Appendix report, you can see how

many of the DL/I calls being issued are distorting the average in the IWAITS/CALL field. See "Distribution-Appendix report" on page 727 for additional ways in which information from the DL/I-Call-Summary report can be used.

Remember, in tuning to decrease I/O operations, pay most attention to calls issued a large number of times. It is more profitable to save 1 second on a call executed 2000 times than to save 5 seconds on a call executed ten times (2000 versus 50 seconds).

## Fields in the DL/I-Call-Summary report

The following example shows a sample DL/I-Call-Summary report.

| IMS MONITOR |             | ****DL/I CALL SUMMARY**** |      |       | TRACE START 1989 076 12:42:54 |              |                | TRACE STOP 1989 076 2:43:07 PAGE 0009 |           |               |                 |  |
|-------------|-------------|---------------------------|------|-------|-------------------------------|--------------|----------------|---------------------------------------|-----------|---------------|-----------------|--|
|             |             | CALL                      | LEV  | STAT  | DL/I                          |              | (C)            | (A)                                   |           | (B)           |                 |  |
| PCB NAME    | CALL FUNC   | NO.SEGMENT                | CODE | CALLS | IWAITS                        | IWAITS/ CALL | ..ELAPSED MEAN | TIME... MAXIMUM                       | .NOT MEAN | IWAIT MAXIMUM | DISTRIB. NUMBER |  |
| DLVNTZ02    | STAT        | (00)                      | GA   | 1     | 1                             | 1.00         | 968281         | 968281                                | 932428    | 932428        | 1 A,B,C         |  |
|             | GN          | (00)                      | GB   | 7     | 0                             | 0.00         | 5703           | 29519                                 | 5703      | 29519         | 4 A,B,C         |  |
|             | GN          | (03)K1Z                   |      | 6     | 0                             | 0.00         | 165            | 253                                   | 165       | 253           | 5 A,B,C         |  |
|             | GN          | (02)K8K5                  | GA   | 8     | 0                             | 0.00         | 263            | 888                                   | 263       | 888           | 6 A,B,C         |  |
|             | GN          | (03)K1Z                   | GK   | 3     | 0                             | 0.00         | 6844           | 20272                                 | 6844      | 20272         | 7 A,B,C         |  |
|             | GN          | (03)K6Y                   |      | 3     | 0                             | 0.00         | 140            | 173                                   | 140       | 173           | 8 A,B,C         |  |
|             | GN          | (03)K5YK1                 | GA   | 3     | 0                             | 0.00         | 197            | 219                                   | 197       | 219           | 9 A,B,C         |  |
|             | GN          | (04)K1Y                   | GK   | 5     | 0                             | 0.00         | 306            | 892                                   | 306       | 892           | 10 A,B,C        |  |
|             | GN          | (04)K42                   |      | 5     | 0                             | 0.00         | 121            | 173                                   | 121       | 173           | 11 A,B,C        |  |
|             | GN          | (03)K5XK2                 | GK   | 5     | 0                             | 0.00         | 9951           | 41900                                 | 9951      | 41900         | 12 A,B,C        |  |
|             | GN          | (03)K6                    |      | 7     | 0                             | 0.00         | 1292           | 7219                                  | 1292      | 7219          | 13 A,B,C        |  |
|             | GN          | (02)K5                    | GA   | 9     | 0                             | 0.00         | 160            | 220                                   | 160       | 220           | 14 A,B,C        |  |
|             | GN          | (04)K1Y                   |      | 7     | 0                             | 0.00         | 12008          | 45991                                 | 12008     | 45991         | 15 A,B,C        |  |
|             | GN          | (03)K5XK2                 |      | 4     | 0                             | 0.00         | 6333           | 20917                                 | 6333      | 20917         | 16 A,B,C        |  |
|             | GN          | (02)K5                    |      | 3     | 0                             | 0.00         | 126            | 138                                   | 126       | 138           | 17 A,B,C        |  |
|             | GN          | (01)K1                    | GA   | 5     | 0                             | 0.00         | 6773           | 23625                                 | 6773      | 23625         | 18 A,B,C        |  |
|             | GN          | (03)K5YK1                 |      | 10    | 0                             | 0.00         | 9444           | 30320                                 | 9444      | 30320         | 19 A,B,C        |  |
|             | GN          | (04)K6X                   | GK   | 4     | 0                             | 0.00         | 141            | 165                                   | 141       | 165           | 20 A,B,C        |  |
|             | GN          | (04)K1X                   | GK   | 4     | 0                             | 0.00         | 3278           | 12578                                 | 3278      | 12578         | 21 A,B,C        |  |
|             | GN          | (04)K4                    |      | 4     | 0                             | 0.00         | 423            | 1342                                  | 423       | 1342          | 22 A,B,C        |  |
|             | GN          | (03)K3K5                  |      | 6     | 0                             | 0.00         | 1360           | 6982                                  | 1360      | 6982          | 23 A,B,C        |  |
|             | GN          | (02)K2                    |      | 4     | 0                             | 0.00         | 325            | 910                                   | 325       | 910           | 24 A,B,C        |  |
|             | GN          | (03)K3K5                  | GA   | 2     | 0                             | 0.00         | 189            | 196                                   | 189       | 196           | 25 A,B,C        |  |
|             | GN          | (04)K1X                   |      | 2     | 0                             | 0.00         | 134            | 142                                   | 134       | 142           | 26 A,B,C        |  |
|             | GU          | (01)K1                    |      | 9     | 0                             | 0.00         | 3387           | 26518                                 | 3387      | 26518         | 27 A,B,C        |  |
|             | GN          | (02)K8K5                  |      | 3     | 0                             | 0.00         | 179            | 207                                   | 179       | 207           | 48 A,B,C        |  |
|             | GN          | (02)K5                    | GE   | 1     | 0                             | 0.00         | 116            | 116                                   | 116       | 116           | 49 A,B,C        |  |
|             | GU          | (03)K5YK1                 |      | 8     | 0                             | 0.00         | 7752           | 26664                                 | 7752      | 26664         | 50 A,B,C        |  |
|             | GNP         | (02)K5                    | GE   | 2     | 0                             | 0.00         | 108            | 120                                   | 108       | 120           | 51 A,B,C        |  |
|             | GNP         | (03)K5YK1                 |      | 5     | 0                             | 0.00         | 7873           | 37362                                 | 7873      | 37362         | 52 A,B,C        |  |
|             | GU          | (04)K1Y                   |      | 1     | 0                             | 0.00         | 669            | 669                                   | 669       | 669           | 56 A,B,C        |  |
|             | GU          | (04)K4                    |      | 4     | 0                             | 0.00         | 11572          | 44675                                 | 11572     | 44675         | 59 A,B,C        |  |
|             | GU          | (02)K5                    |      | 6     | 0                             | 0.00         | 9025           | 45363                                 | 9025      | 45363         | 61 A,B,C        |  |
|             | GNP         | (02)K2                    | GE   | 1     | 0                             | 0.00         | 87             | 87                                    | 87        | 87            | 62 A,B,C        |  |
|             | GNP         | (03)K3K5                  |      | 2     | 0                             | 0.00         | 165            | 167                                   | 165       | 167           | 63 A,B,C        |  |
|             | GU          | (02)K2                    |      | 3     | 0                             | 0.00         | 431            | 566                                   | 431       | 566           | 64 A,B,C        |  |
|             | GU          | (03)K5XK2                 |      | 5     | 0                             | 0.00         | 13570          | 52617                                 | 13570     | 52617         | 65 A,B,C        |  |
|             | GU          | (04)K6X                   |      | 3     | 0                             | 0.00         | 4892           | 13242                                 | 4892      | 13242         | 66 A,B,C        |  |
|             | GU          | (03)K6                    |      | 2     | 0                             | 0.00         | 1247           | 1810                                  | 1247      | 1810          | 67 A,B,C        |  |
|             | GU          | (04)K42                   |      | 3     | 0                             | 0.00         | 998            | 1390                                  | 998       | 1390          | 68 A,B,C        |  |
|             | DLET        | (04)K42                   |      | 2     | 0                             | 0.00         | 46374          | 49018                                 | 46374     | 49018         | 72 A,B,C        |  |
|             | DLET        | (04)K6X                   |      | 2     | 0                             | 0.00         | 108321         | 118802                                | 108321    | 118802        | 73 A,B,C        |  |
|             | DLET        | (02)K2                    |      | 1     | 0                             | 0.00         | 44789          | 44789                                 | 44789     | 44789         | 74 A,B,C        |  |
|             | DLET        | (03)K3K5                  |      | 3     | 0                             | 0.00         | 87567          | 121071                                | 87567     | 121071        | 75 A,B,C        |  |
|             | GU          | (02)J5                    | GE   | 1     | 0                             | 0.00         | 1644           | 1644                                  | 1644      | 1644          | 111 A,B,C       |  |
|             | DLET        | (03)J6                    |      | 1     | 0                             | 0.00         | 35633          | 35633                                 | 35633     | 35633         | 112 A,B,C       |  |
|             | GU          | (02)J9                    | GE   | 1     | 0                             | 0.00         | 705            | 705                                   | 705       | 705           | 113 A,B,C       |  |
|             | DLET        | (04)J7J9                  |      | 1     | 0                             | 0.00         | 245921         | 245921                                | 245921    | 245921        | 114 A,B,C       |  |
|             | REPL        | (03)J7PJ6                 | DA   | 1     | 0                             | 0.00         | 225            | 225                                   | 225       | 225           | 115 A,B,C       |  |
|             | BATCH TOTAL |                           |      |       |                               |              |                |                                       |           |               |                 |  |
|             |             |                           |      | 392   | 22                            | 0.05         | 22310          |                                       | 20983     |               |                 |  |

The meaning of the various fields in the report is as follows:

### TRACE START and TRACE STOP

The trace start and trace stop fields tell you the time when the DB Monitor program was last started and stopped. The time is generated by the time-of-day-clock. Clock times are read as follows:

Clock time = hh.mm.ss

where:

**hh**      Hours 0 through 23

**mm**      Minutes

**ss**      Seconds

If the DB Monitor was on during an entire batch run, the trace start and trace stop time is when the batch run started and stopped. If the DB Monitor was turned on and off more than once in the same batch run, the trace start and trace stop times are when the monitor was last started and stopped.

**PCBNAME**

This field describes the name of the PCB the report is providing information about. Remember that each application program has one or more PCBs. If PCB names are unique to application programs, this field can only be used to identify which application program the report is providing information about.

**CALL FUNC**

This field describes the type of DL/I call the application program issued.

**LEV NO.**

This field describes the level that was accessed in the hierarchy of the database record to perform the DL/I call. If this field contain zeros, it means position was not established and therefore no level was set. This generally happens when a DL/I call cannot be processed for some reason.

**SEGMENT**

This field describes the 8-character name of the segment accessed by the DL/I call.

**STAT CODE**

This field describes the status code returned after the call (if the status code was not blank).

**DL/I Calls**

This field describes how many times this particular DL/I call was issued and had the five unique characteristics listed in the previous five columns.

**IWAITS**

This field describes the number of times IMS was inactive.

**IWAITS/CALL**

This field describes, by DL/I call, the average number of times IMS was inactive.

**ELAPSED TIME**

This field describes, by DL/I call, the elapsed time for calls.

**MEAN**

The average elapsed time per DL/I call

**MAXIMUM**

The longest elapsed time for a single DL/I call

**NOT IWAIT TIME**

This field describes, by DL/I call, the elapsed time minus the IWAIT time.

**MEAN**

The average NOT IWAIT TIME

**MAXIMUM**

The longest single NOT IWAIT TIME

NOT IWAIT TIME includes any time spent by higher priority tasks running in the IMS region. NOT IWAIT TIME might be about equal to total processor time if the IMS database region is the high priority task and no low priority tasks are causing interrupts.

**DISTRIB. NUMBER**

The distribution number is a reference number to be used in conjunction with the Distribution-Appendix report. For a description of what it means, see "Distribution-Appendix report."

**C, A, and B**

These letters over the IWAITS/CALL, ELAPSED TIME, and NOT IWAIT TIME columns are reference letters to be used in conjunction with the Distribution-Appendix report. For a description, see "Distribution-Appendix report."

**BATCH Total**

For this execution of the DB Monitor this field describes:

**DL/I Calls**

The total number of DL/I calls

**IWAITS**

The total number of times IMS was inactive

**IWAITS/CALL**

The average number of IWAITS per DL/I call

---

**Distribution-Appendix report**

Use a Distribution-Appendix report when you suspect some unusual combination of events has occurred, and the times or totals on the Program-I/O or DL/I-Call-Summary report do not give you enough information to highlight the problem.

The Distribution-Appendix report takes specific events for which a time or a total has been generated and distributes the events across ranges. It does this for specific events from the Program-I/O report and the DL/I-Call-Summary report.

To get an idea of when the Distribution-Appendix report is useful, see the following example. Suppose in a DL/I-Call-Summary report, the row of information in the example appears:

| PCB     | CALL | LEV NO.      | DL/I  | (C)    |      |          |  |
|---------|------|--------------|-------|--------|------|----------|--|
| NAME    | FUNC | SEGMENT      | CALLS | IWAITS | ...  | DISTRIB. |  |
| ----    | ---- | -----        | ----- | CALL   | ---- | NUMBER   |  |
| DHVB203 | DLET | (01)A1111111 | 11    | 8.63   |      | 11C      |  |

The (C) over the IWAITS/CALL field means that this event is detailed on the Distribution-Appendix report. The DISTRIB. NUMBER field says this event is broken down on the Distribution-Appendix report specifically on line 11 C. In this example, the IWAITS/CALL field value of 8.63 is a very high number relative to the other IWAITS/CALL numbers on the DL/I-Call-Summary report. For more information, check line 11 C in the Distribution-Appendix report. It looks like this:

11 C ....0....0....1....2....3....4....5....6....7....8....INF  
          0      0      0      0      0      2      4      1      3      1

The numbers next to line 11 C (0 to INF) are predefined ranges. The numbers beneath line 11 C are the distribution of the event. Look for more information

about why the IWAITS/CALL total is high (for DLET calls issued under PCB DHVBT203 against segment A111111). The distribution of numbers beneath line 11 C says that of the 11 DLET calls issued:

- Two calls caused 5 IWAITS
- Four caused 6 IWAITS
- One caused 7 IWAITS
- Three caused 8 IWAITS
- One caused more than 8 IWAITS

Because one call caused more than 8 IWAITS, this call might be distorting the average in the IWAITS/CALL field in the report. Tune your database to eliminate relatively high IWAITS per DL/I call, investigate the call that required more than 8 IWAITS. In this example, the interval between distributions is 1. For other entries (for example, line 4B in the report in the following example), the intervals are much larger than 1. In these cases, interpret the data as (using line 4B): 1 call fell in the range of 16000 to 32000.

The example is of a Distribution-Appendix report. To read it, remember:

- The lines in the Distribution-Appendix report are always identified by a number (55, 56, and so on) or a number and a character (3A, 3B, and so on). The numbers are the numbers used in the DISTRIBUTION NUMBER fields in the Program-I/O or DL/I-Call-Summary report. The characters are used in the DL/I-Call-Summary report to identify which event is being distributed.
- The numbers next to the line are predefined ranges across which an event can be distributed. These ranges are made appropriate to the event. For example, ranges 0, 1, 2, 3, and so on are appropriate for distributing IWAIT totals for DL/I calls. (In the example used, we wanted to know how many calls required 0, 1, 2, 3, and so on IWAITS.) Ranges such as 0, 1000, 2000, and so on are appropriate for all other events in the Program-I/O and DL/I-Call Summary reports because all other events are times, not totals. These 0, 1000, 2000, and so on numbers are in microseconds. A microsecond is one millionth of a second (in other words, 2000 microseconds equals 0.002000 second). The ranges have default values, but you can redefine the ranges. See “Redefining the default ranges” on page 730 for more information.
- The number below the line is always the number for the event being distributed.

```
IMS MONITOR  ****DISTRIBUTION APPENDIX****  RACE START 1989 076 12:42:54  TRACE STOP 1989 076 12:43:07 PAGE 0014

# 1A.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
    0          0          0          0          0          0          0          0          0          0          0          1
1B.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
    0          0          0          0          0          0          0          0          0          0          0          1
1C.....0.....0.....1.....2.....3.....4.....5.....6.....7.....8.....INF
    0          1          0          0          0          0          0          0          0          0          0          0
# 4A.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
    5          0          0          1          0          0          1          0          0          0          0          0
4B.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
    5          0          0          1          0          0          1          0          0          0          0          0
4C.....0.....0.....1.....2.....3.....4.....5.....6.....7.....8.....INF
    7          0          0          0          0          0          0          0          0          0          0          0
# 5A.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
    6          0          0          0          0          0          0          0          0          0          0          0
5B.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
    6          0          0          0          0          0          0          0          0          0          0          0
5C.....0.....0.....1.....2.....3.....4.....5.....6.....7.....8.....INF
    6          0          0          0          0          0          0          0          0          0          0          0
# 6A.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
    8          0          0          0          0          0          0          0          0          0          0          0
6B.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
    8          0          0          0          0          0          0          0          0          0          0          0
6C.....0.....0.....1.....2.....3.....4.....5.....6.....7.....8.....INF
    8          0          0          0          0          0          0          0          0          0          0          0
```

```

# 7A.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
      2          0          0          0          0          1          0          0          0          0
7B.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
      2          0          0          0          0          1          0          0          0          0
7C.....0.....0.....1.....2.....3.....4.....5.....6.....7.....8.....INF
      3          0          0          0          0          0          0          0          0          0
# 8A.....0.....1000.....2000.....4000.....8000.....16000.....32000.....64000.....128000.....256000.....INF
      3          0          0          0          0          0          0          0          0          0
:
# 131.....0.....2000.....8000.....24000.....50000.....100000.....150000.....200000.....250000.....300000.....INF
      0          1          5          2          0          0          0          0          0          0
# 132.....0.....2000.....8000.....24000.....50000.....100000.....150000.....200000.....250000.....300000.....INF
      0          0          5          1          0          0          0          0          0          0
# 133.....0.....2000.....8000.....24000.....50000.....100000.....150000.....200000.....250000.....300000.....INF
      0          0          3          0          0          0          0          0          0          0

```

The TRACE START and TRACE STOP fields at the top of the report tell you the time when the DB Monitor was last started and stopped. The time is generated by the time-of-day clock. Clock times are read as follows:

Clock time = hh.mm.ss

Where:

**hh** Hours 0 through 23

**mm** Minutes

**ss** Seconds

If the DB Monitor was on during an entire batch run, the trace start and trace stop times are the times at which the batch run started and stopped. If the DB Monitor was turned on and off more than once in the same batch run, the trace start and trace stop times are the times at which the monitor was last started and stopped.

## How to generate the Distribution-Appendix report

The Distribution-Appendix report is not generated automatically when the DB Monitor is run. To get the Distribution-Appendix report, you have to include a DIS input control statement in the analysis control data set.

## Events that can be distributed and their default ranges

Certain events, such as the IWAIT time for OSAM, can be distributed in the Program-I/O and DL/I-Call-Summary reports.

The following table shows the events that can be distributed in the Program-I/O and DL/I-Call-Summary reports. Each event has an ID.

*Table 56. Events that can be distributed and their IDs.*

| Events that can be distributed | ID  |
|--------------------------------|-----|
| Program-I/O report             |     |
| IWAIT time for OSAM            | D23 |
| IWAIT time for VSAM            | D24 |
| IWAIT time for HSAM            | D34 |
| Elapsed time per DL/I call     | D11 |
| Not IWAIT time per DL/I call   | D12 |
| IWAITs per DL/I call           | D13 |



The following table shows, using this ID, what each predefined set of ranges consists of when the default ranges are used. Notice that the first number in a range always defaults to zero, and the last number always defaults to infinity (INF).

*Table 57. Predefined ranges when the default is used*

| ID      | Default ranges                                                 |
|---------|----------------------------------------------------------------|
| D11,D12 | 0,1000,2000,4000,8000,16000,32000,64000,128000,256000,INF      |
| D13     | 0,0,1,2,3,4,5,6,7,8,INF                                        |
| D23,D24 | 0,2000,8000,24000,50000,100000,150000,200000,250000,300000,INF |
| D34     | 0,2000,4000,8000,16000,32000,64000,96000,128000,160000,INF     |

## Redefining the default ranges

You can redefine the default ranges. To do this, you have to include an input control statement in the analysis control data set for each default range you want to override.

**Related reading:** The analysis control data set is explained under “IMS Monitor Report Print utility (DFSUTR20)” in *IMS Version 11 Database Utilities*.

To override default ranges, you specify control statements in the form Dn n1,n2,...where:

- Dn is the ID of the event to be distributed (D23, for example, is the ID for IWAIT time for the OSAM event in the Program-I/O report)
- n1 is a specific default value. Each of the 10 default values or buckets can be redefined. For example, the Program-I/O report IWAIT time for OSAM could be redefined as follows:

```
D23 0,500,1000,1500,2000,4000,,,100000,500000
```

This would result in the 7th and 8th ranges remaining at their default values (150000 and 200000) and the final range (INF) remaining at its default value.

## Monitor-Overhead report

This report describes the overhead required to run the DB Monitor. Because the DB monitor runs while IMS is executing, the monitor's use of resources causes the performance to be slightly less than it is when the monitor is not on.

### Fields in the Monitor-Overhead report

The following example shows a sample of a Monitor-Overhead report.

```
IMS MONITOR  ****MONITOR OVERHEAD****      TRACE START 1989 076  12:42:54      TRACE STOP 1989 076  12:43:07 PAGE 0013
MONITOR OVERHEAD DATA                                     M
                                                           M
13144 MILLISECONDS, TRACE INTERVAL
550 MILLISECONDS, MONITOR MODULE TIME
853 MONITOR RECORDS WERE PRODUCED
645 MICROSECONDS PER MONITOR ENTRY
```

The meaning of the various fields in the report is as follows:

#### TRACE START and TRACE STOP

The trace start and trace stop fields tell you the time when the DB Monitor program was last started and stopped. The time is generated by the time-of-day clock. Clock times are read as follows:

Clock time = hh.mm.ss

Where:

**hh**      Hours 0 through 23

**mm**      Minutes

**ss**      Seconds

If the DB Monitor was on during an entire batch run, the trace start and trace stop times are the times at which the batch run started and stopped. If the DB Monitor was turned on and off more than once in the same batch run, the trace start and trace stop times are the times at which the monitor was last started and stopped.

**MILLISECONDS, TRACE INTERVAL**

This field describes how long the monitor was turned on.

**MILLISECONDS, MONITOR MODULE TIME**

This field describes, during the time the monitor was turned on, how long code in the monitor was actually being executed. (During this time, no IMS code can be executed, so this field tells the real overhead for using the monitor.)

**MONITOR RECORDS WERE PRODUCED**

This field describes how many records the monitor wrote.

**MICROSECONDS PER MONITOR ENTRY**

This field describes the average length of time it took the monitor to write a record.



---

## Chapter 48. IMS Monitor reports

The IMS Monitor output reflects the IMS DB/DC environment as a whole. A subset of the reports deals with database calls and buffering.

Interpreting this data for batch application performance and to verify database design are considered separate tasks allied with database administration.

---

### Transaction flow and IMS Monitor events

For an overall picture of the events, system activities, and usage of storage areas (buffer pool or data set) for which the IMS Monitor gathers timings, see the following table.

The leftmost column shows, from top to bottom, a sequence of processing events; each event is related to IMS Monitor reported items in the notes that follow the table.

*Table 58. Transaction flow and IMS Monitor events description.*

| Flow | Event                    | Activity                  | Pool               | Data set      |
|------|--------------------------|---------------------------|--------------------|---------------|
| 1    | Wait for poll            | Waiting                   | N/A                | N/A           |
| 2    | Data Transfer            | N/A                       | N/A                | N/A           |
| 3    | Input message processing | Device                    | CIOP               | (For lines)   |
|      |                          | MFS                       | MFP                | IMS.FORMATx   |
|      |                          | Enqueuing                 | QBUF               | Message queue |
| 4    | Input queuing            | Waiting                   | N/A                | N/A           |
| 5    | Scheduling               | Scheduling                | QBUF               | Message queue |
|      |                          | PSB load                  | PSB, PSBW,<br>DPSB | IMS.ACBLIBx   |
|      |                          | DMB load                  | DMBP               | IMS.ACBLIBx   |
| 6    | Program load             | Program load              |                    | IMS.PGMLIB    |
| 7    | Program initialization   | Initializing              | N/A                | N/A           |
| 8    | Message queue GU         | Primed GU call            | N/A                | N/A           |
| 9    | Program Execution        | DC calls                  | QBUF               | Message queue |
|      | DL/I Elapsed             | DB calls                  | OSAM/VSAM          | N/A           |
|      | Wait Elapsed             | DB I/Os                   | OSAM/VSAM          | Databases     |
|      |                          | SPA insert                | CWAP               | IMS.SPA       |
| 10   | Output message insert    | DC ISRT call              | QBUF               | Message queue |
|      |                          | DC GU call                | QBUF               | Message queue |
|      |                          | Sync point<br>(MODE=SNGL) | OSAM/VSAM          | Databases     |
| 11   | Wait for sync point      | (Only for<br>MODE=MULT)   | N/A                | N/A           |
| 12   | Program termination      | Program<br>termination    | OSAM/VSAM          | Databases     |

Table 58. Transaction flow and IMS Monitor events description (continued).

| Flow | Event                     | Activity     | Pool | Data set      |
|------|---------------------------|--------------|------|---------------|
| 13   | Wait for selection        | Waiting      | QBUF | Message Queue |
| 14   | Output message processing | Message send | QBUF | Message queue |
|      |                           | MFS          | MFP  | IMS.FORMATx   |
|      |                           | Device       | CIOP | (For lines)   |
| 15   | Data transfer             | N/A          | N/A  | N/A           |
| 16   | Output queue processing   | Dequeuing    | QBUF | Message queue |

**Notes:**

1. The time waiting for poll is not recorded by the IMS Monitor.
2. Line activity in terms of data transmitted and IMS communication sub-task activity are recorded for all messages. Input message activity is not separated from output message activity.
3. During input message processing, the device dependent processing and the use of the communication I/O pool (CIOP) are recorded. If Message Format Service (MFS) is required, the use of the message format buffer pool and any I/O to the active IMS.FORMATA/B library are recorded. The input message is then placed in a message queue buffer (QBUF) with possible I/O to the message queue data sets. The IMS Monitor reporting does not distinguish this activity from output message processing.
4. Time spent waiting on the input queue is not recorded by the IMS Monitor.
5. IMS schedules the processing program into a region and, as part of this action, accesses the message queue or the QBUF pool so that it can present the message to the program. At this time, the required PSB and physical database blocks are made available in the PSB pool and the DMB pool. If they are not already in the pool, they are retrieved from the active IMS.ACBLIBA/B data set. These events are summarized under SCHEDULING AND TERMINATION in IMS Monitor reports.
6. Next, the application program is loaded into the region from IMS.PGMLIB, or initialized if already resident in the region. This processing is included as part of SCHEDULE TO FIRST CALL in IMS Monitor reports.
7. The initialization performed by the application program is included in this period of time, up to the time of the first message queue or database call. This processing is considered part of the scheduling process because it is not repeated when multiple transactions are processed for one scheduling event. This processing is recorded as part of the SCHEDULE TO FIRST CALL in IMS Monitor reports.
8. The event of performing the first DL/I GU call to obtain the first message segment is not separately recorded. The processing to prime the application with this message is included with the SCHEDULE TO FIRST CALL.
9. The program elapsed event is measured from the first call to termination of the processing program. The total elapsed time is recorded in IMS Monitor reports as ELAPSED EXECUTION.
  - Each DL/I call, whether DC or DB, is individually recorded along with its use of message queues or database data sets and their respective pools. Each external subsystem call is individually recorded. These events are recorded under CALLS in IMS Monitor reports.
  - When a DL/I call causes I/O activity, the time spent waiting for the data is recorded as WAIT time and as the number of I/Os. When a program's

database processing intent and other update activity are in contention, this also contributes to WAIT time. For each external subsystem call, the time spent in external subsystem processing is recorded separately as WAIT time. When processing is suspended because of intent, the elapsed time is recorded in IDLE FOR INTENT IMS Monitor report items. When processing is suspended because the region is designated as wait-for-input and no input message is available, the time spent waiting for the next input message is excluded from elapsed time intervals and wait times. Wait-for-input time appears only in \*\*WFI items on the program I/O report.

- If the message processing is part of a conversational transaction, the activity in the communications work area pool (CWAP) and the IMS.SPA data set is recorded.
10. The DL/I ISRT call for the response to the message just processed is recorded for the transaction.  
As processing continues, there are synchronization points such as a GU to the message queue for another input message or a checkpoint call. The database and message queue I/O to commit the program processing is recorded. Checkpointing by IMS or by the processing program is recorded separately.
  11. If the processing is for multiple messages before synchronization (MODE=MULT), there can be a wait for sync point. This time is not recorded by the IMS Monitor.
  12. The processing after exiting from the application program, the program termination events, is included with the initial scheduling time. It is part of SCHEDULING AND TERMINATION.
  13. The output message waits for the selection to be transmitted to the terminal. The duration of any wait time on the output queue is not recorded by the IMS Monitor.
  14. After the exit from the application program and termination, output message processing events are recorded. These include message format and device dependent processing as well as sending the output message. The IMS Monitor reporting does not distinguish this activity from input message processing.
  15. Line activity in terms of data transmitted and IMS communication sub-task activity are recorded for all messages. Output message activity is not separated from input message activity.
  16. The processing to remove the output message from the queue is recorded.

---

## IMS Monitor trace event intervals

The IMS Monitor trace interval is bounded by the master terminal operator's use of the /TRACE command between the start and stop command entries. The IMS Monitor can also be stopped by the expiration of any interval specified.

The online IMS events are recorded in IMS Monitor records placed in the IMSMON data set. The event timings are related to dependent region activity. The following figure shows the boundaries of the timed event intervals.

The Monitor trace interval includes the following intervals:

- Scheduling and Termination
  - No Messages
  - Block loader busy
  - Intent failures (exclusive intent and data sharing) and Schedule failures (PSB busy and space failure)

- Sched/Term elapsed
  - NOT-WAIT
  - ACBLIB waits
  - DB Flush waits
  - DB CLOSE waits
- Region occupancy (which overlaps with all of Sched/Term elapsed)
  - Schedule to first call
  - Elapsed execution

The NOT-WAIT time for a region is the elapsed time not accounted for by wait time. Any delay coming from either paging or the processor being dispatched for a higher priority task results in an increase in the NOT-WAIT times.




Figure 87. IMS Monitor trace event intervals

## Overview of IMS Monitor reports

A list of reports available from data collected by the IMS Monitor, together with the principal performance data they contain, is shown in the following table.

**Related reading:** For those reports marked by a “DC”, see Chapter 50, “IMS Monitor reports for DCCTL,” on page 787.

The reports marked “MSC” in the list are only produced when MSC is active. The MSC reports and their interpretation are in the following table.

The order of the reports listed in the following table matches the sequence of the output from the IMS Monitor Report Print Program. The duration or constraints of a monitoring snapshot might not include certain events necessary for an individual report, in which case only report headings or partial data are produced.

Table 59. IMS Monitor reports output sequence and information

| Report name          | Principal information                     |
|----------------------|-------------------------------------------|
| System Configuration | Monitor run documentation                 |
| Message Queue Pool   | Buffering and message I/O per transaction |

*Table 59. IMS Monitor reports output sequence and information (continued)*

| Report name                | Principal information                     |
|----------------------------|-------------------------------------------|
| Database Buffer Pool (DB)  | Count of DB calls and I/O per transaction |
| VSAM Buffer Pool (DB)      | Count of inserts and I/Os                 |
| Message Format Buffer Pool | Count of blocks fetched and I/Os          |
| Latch Conflict Statistics  | IMS internal processing                   |
| General Wait Time Events   | Wait times for SNAPQ                      |
| Region and Jobname         | Monitor run documentation                 |
| Region Summary             | Elapsed times and count of DL/I calls     |
| Region Wait                | Wait times                                |
| Programs by Region         | Elapsed times for region usage            |
| Program Summary            | Overall program statistics                |
| Program I/O (DB)           | Wait times/PCB                            |
| Communication Summary      | Elapsed times for lines                   |
| Communication Wait         | Wait times by line                        |
| Line Functions             | Count and size of blocks transmitted      |
| MSC Traffic (MSC)          | Count and routing of transactions         |
| MSC Summaries (MSC)        | Count of transactions by destination      |
| MSC Queuing Summary (MSC)  | Count and queuing time by link            |
| Transaction Queuing        | Queue loading statistics                  |
| Reports                    | Count of space failures and deadlocks     |
| Run Profile                | Monitor run documentation                 |
| Call Summary (DB)          | Call counts and timings/segment type      |
| Distribution Appendix      | Event frequency distributions             |

The majority of the data items in IMS Monitor reports are elapsed times. These are normally expressed in microseconds. An entry of 1876534 represents 1.876534 seconds or 1876 milliseconds. Any times that do not follow this convention show the unit of measure on the report.

You can also find counts of events under the heading OCCURRENCES, and some figures that represent the number of bytes.

## Documenting the monitoring run

For each trace interval there are several general reports or overall summaries of the processing that took place. You can use these reports as part of your IMS Monitor run documentation.

It is important to record, as accurately as possible the conditions under which the trace was taken. Your documentation can include system status information obtained by the /DISPLAY command several times before and after the trace, an expected profile of the application program activity, and any desired processing events. The trace interval should represent typical processing loads and not be a biased or inadequate historical record.



## Adding to the System-Configuration report data

The first general report titled SYSTEM CONFIGURATION is found under the page heading BUFFER POOL STATISTICS. It shows the modification level of the IMS and z/OS systems. The system configuration output is illustrated in the following section.

## Recording the Monitor trace interval

The heading of most IMS Monitor reports carries the trace start and stop times. It is shown in the format YEAR DAY (Julian) HH:MM:SS. The overall length of the trace interval is given in milliseconds under the title MONITOR OVERHEAD DATA. The following line shows how many trace records were placed on the IMSMON data set. An example of the monitor trace interval recording is shown in the following example.

```
***IMS MONITOR***  BUFFER POOL STATISTICS  TRACE START 2009 180   5:55:15   TRACE STOP  2009 180   5:59:49   PAGE 0001
S Y S T E M   C O N F I G U R A T I O N

                                SYSTEM CONFIGURATION      :  MVS/ESA
                                IMS VERSION                 :  11
                                RELEASE LEVEL                :   1
                                MODIFICATION NUMBER          :   0
```

## Completing the Monitor run profile

A compact set of processing ratios is found at the end of the Run-Profile report. The statistics summarize, for the monitor interval, the transaction throughput and the degree of DL/I and I/O activity. An example of the run-profile report is shown in the following example.

The lower part of the Run-Profile report shows several ratios:

- Program elapsed time to DL/I elapsed time for each region
- DL/I elapsed time to wait time during DL/I processing
- Program elapsed time to other subsystem call elapsed time
- DL/I elapsed time to other subsystem call elapsed time

Each dependent region is identified by a sequence number, starting at region 1.

```
IMS MONITOR  **RUN PROFILE**      TRACE START 2009 180   5:55:15   TRACE STOP  2009 180   5:59:49   PAGE 0184
TRACE ELAPSED TIME IN SECONDS.....274.6
TOTAL NUMBER OF MESSAGES DEQUEUED.....1403
TOTAL NUMBER OF SCHEDULES.....173
NUMBER OF TRANSACTIONS PER SECOND.....5.1
TOTAL NUMBER OF DL/I CALLS ISSUED.....18632
NUMBER OF DL/I CALLS PER TRANSACTION.....13.2
NUMBER OF OSAM BUFFER POOL I/O'S.....11236,      8.0 PER TRANSACTION
NUMBER OF MESSAGE QUEUE POOL I/O'S.....0,         0.0 PER TRANSACTION
NUMBER OF FORMAT BUFFER POOL I/O'S.....0,         0.0 PER TRANSACTION
RATIO OF PROGRAM ELAPSED TO DL/I ELAPSED:
    REGION  1:   1.09
    REGION  2:   1.09
    REGION  3:   1.00
    REGION  4:   1.02
    REGION  5:   1.01
    REGION  6:   1.00
    REGION  7:   1.00
    REGION  8:   1.00
    REGION  9:   1.17
    REGION 10:   1.00
    REGION 11:   1.00
    REGION 49:   1.03
    REGION 50:   1.19
RATIO OF DL/I ELAPSED TO DL/I IWAIT:
    REGION  1:  325.65
```

```

REGION 2: 73.49
REGION 4: 100.35
REGION 5: 85.76
REGION 6: 82.99
REGION 47: 95.64
REGION 48: 45.93
REGION 49: 9.22

```

You can match the regions to the name of the z/OS job using the Region-and-Jobname report. The job names correspond to the step names on the EXEC statements of all the dependent regions started by the operator before the trace was started. The region job names are included on the monitor output page with the heading GENERAL REPORTS, as illustrated in “Detecting checkpoint effects” on page 751.

Some generalized processing ratios are given at the end of several buffer pool statistics reports. You can include them in the documented profile of the trace interval. These are not specific to one application or system resource but can be used as indicators of variation across a series of monitor runs.

The ratios are:

- The total number of OSAM reads + OSAM writes + all waits divided by the total number of transactions.  
From the Message-Queue-Pool report, this ratio indicates on a per transaction basis the physical I/O activity required to handle the message queuing function.
- The total number of OSAM reads + OSAM writes + BISAM reads divided by the total number of transactions.  
From the Database-Buffer-Pool report, this ratio indicates on a per transaction basis the physical I/O activity required to handle the database buffering function.
- The total prefetch I/Os + immediate fetch I/Os + directory I/Os divided by the total number of transactions.  
From the Message Format Buffer Pool report, this ratio indicates on a per transaction basis the physical I/O activity required to handle the MFS function.

## Verifying IMS Monitor report occurrences

When you examine the output from the IMS Monitor Report Print program, the presence of a report heading does not necessarily mean that appropriate data is listed.

System definition options and utility control statements affect the content of the output as follows:

- The output does not include a Call-Summary report unless a control statement specifies DLI.
- The output does not include a set of Distribution reports unless a control statement specifies DIS or DISTRIBUTION. The column headed DISTRIBUTION NUMBER that occurs on many of the reports contains cross-references to items included in the Distribution reports.
- The output consists of just a Call Summary report if a control statement specifies ONLY DLI.

Because many of the summary reports require system status to calculate the difference between start and end values, and this status is obtained during the

/TRACE SET OFF processing, the IMS Monitor execution must end before termination of the IMS control region. If the trace was not stopped properly, the following message is issued:

```
NO QUEUE BUFFER POOL TRACES AT END TIME ON MONITOR LOG TAPE
****QUEUE BUFFER POOL REPORT CANCELLED****
```

Similarly, other summary reports are not produced.

The series of reports with the title BUFFER POOL STATISTICS do not include a VSAM BUFFER POOL section unless one of the databases in IMS.ACBLIB uses the VSAM access method. If VSAM is not used, the following message is issued:

```
NO VSAM BUFFER POOL TRACES ON MONITOR LOG TAPE
****VSAM BUFFER POOL REPORT CANCELLED****
```

The section MESSAGE FORMAT BUFFER POOL is included only if your system definition specifies devices using Message Format Service (MFS).

If the source data used to formulate a particular IMS Monitor report, or a section of that report, has not been recorded by the IMS Monitor during the trace interval, the report contains only the headings.

---

## Monitoring activity in dependent regions

The IMS Monitor gathers timing information for every dependent region identified in the /TRACE command active during the trace interval. It records the total of the elapsed times for each event, the maximum individual time encountered, and the average time.

There are three major reports that display timings. The reports and a list of their content are:

- **Region-Summary Report**
  - Scheduling and termination
  - Schedule end to first call
  - Elapsed execution with separate summaries shown for:
    - DL/I calls
    - External subsystem service and command calls
    - External subsystem database access calls
    - Checkpoint processing
    - Region occupancy
- **Region Wait**
  - Waits during scheduling and termination
  - Waits during DL/I calls
  - Waits during external subsystem calls
  - Waits during checkpoint
- **Programs by Region**
  - Elapsed execution
  - Schedule end to first call

These three reports are illustrated in the following examples.

Activities for dependent regions are placed in five categories:

- Elapsed time for scheduling and termination  
The scheduling process includes many preparatory events such as block loading from an active IMS.ACBLIBA/B data set, and obtaining ownership of the PSB. The time required to terminate the region activity after the application program ends is also included.
- Elapsed time from end of schedule to first call  
This time is reserved for application program initialization and housekeeping prior to an initial call (to the message queue, a database, or an external subsystem) that marks the beginning of control program services. It is a measure of processing that is not repeated when multiple transactions are processed in a single scheduling.
- Program elapsed time, including all calls  
This time encompasses the major application program processing, measured from the first call to the return or exit from the program.
- Elapsed time performing DL/I calls  
This time includes all DL/I calls. Each DL/I call event is measured from the time of the call to the return to the application program.
- Elapsed time performing external subsystem calls  
This time includes all external subsystem calls. Each external subsystem event is measured from the time of the call to the return to IMS.

The following example shows a sample of the region-summary report.

| IMS MONITOR                |    | ****REGION SUMMARY**** |             |           |           | TRACE START 1993 130 |         | 5:55:15             | TRACE STOP 1993 130 |             | 5:59:49      | PAGE 0011 |
|----------------------------|----|------------------------|-------------|-----------|-----------|----------------------|---------|---------------------|---------------------|-------------|--------------|-----------|
|                            |    | (A)                    |             |           |           |                      |         | (B)                 |                     |             |              |           |
|                            |    | .....ELAPSED TIME..... |             |           |           | NOT IWAIT            |         | TIME(ELAPSED-IWAIT) |                     |             | DISTRIBUTION |           |
|                            |    | OCCURRENCES            | TOTAL       | MEAN      | MAXIMUM   | TOTAL                | MEAN    | MAXIMUM             |                     |             | NUMBER       |           |
| SCHEDULING AND TERMINATION |    |                        |             |           |           |                      |         |                     |                     |             |              |           |
| **REGION                   | 5  | 5                      | 4146        | 829       | 948       | 4146                 | 829     | 948                 |                     |             | 287A,B       |           |
| **REGION                   | 6  | 7                      | 6028        | 861       | 1067      | 6028                 | 861     | 1067                |                     |             | 214A,B       |           |
| **REGION                   | 8  | 8                      | 6847        | 855       | 1098      | 6847                 | 855     | 1098                |                     |             | 129A,B       |           |
| **REGION                   | 10 | 7                      | 9664        | 1380      | 3668      | 9664                 | 1380    | 3668                |                     |             | 272A,B       |           |
| **REGION                   | 47 | 6                      | 5482        | 913       | 1021      | 5482                 | 913     | 1021                |                     |             | 145A,B       |           |
| **REGION                   | 49 | 3                      | 2612        | 870       | 917       | 2612                 | 870     | 917                 |                     |             | 443A,B       |           |
| **TOTALS                   |    | 123                    | 126042      | 1024      |           | 126042               | 1024    |                     |                     |             |              |           |
| SCHEDULE TO FIRST CALL     |    |                        |             |           |           |                      |         |                     |                     |             |              |           |
| **REGION                   | 1  | 1                      | 15479797    | 15479797  | 15479797  |                      |         |                     |                     |             |              | 555       |
| **REGION                   | 2  | 1                      | 22376350    | 22376350  | 22376350  |                      |         |                     |                     |             |              | 564       |
| **REGION                   | 3  | 1                      | 15169488    | 15169488  | 15169488  |                      |         |                     |                     |             |              | 578       |
| **REGION                   | 4  | 1                      | 48146258    | 48146258  | 48146258  |                      |         |                     |                     |             |              | 584       |
| **REGION                   | 48 | 1                      | 795351      | 795351    | 795351    |                      |         |                     |                     |             |              | 592       |
| **REGION                   | 49 | 4                      | 2960425     | 740106    | 2951746   |                      |         |                     |                     |             |              | 442       |
| **REGION                   | 50 | 1                      | 15713464    | 15713464  | 15713464  |                      |         |                     |                     |             |              | 575       |
| **TOTALS                   |    | 168                    | 514286738   | 3061230   |           |                      |         |                     |                     |             |              |           |
| ELAPSED EXECUTION          |    |                        |             |           |           |                      |         |                     |                     |             |              |           |
| **REGION                   | 1  | 1                      | 290146255   | 290146255 | 290146255 |                      |         |                     |                     |             |              | 1         |
| **REGION                   | 2  | 1                      | 252290108   | 252290108 | 252290108 |                      |         |                     |                     |             |              | 2         |
| **REGION                   | 3  | 1                      | 259496970   | 259496970 | 259496970 |                      |         |                     |                     |             |              | 3         |
| **REGION                   | 4  | 1                      | 322812716   | 322812716 | 322812716 |                      |         |                     |                     |             |              | 4         |
| **REGION                   | 48 | 1                      | 273871107   | 273871107 | 273871107 |                      |         |                     |                     |             |              | 48        |
| **REGION                   | 49 | 4                      | 271703421   | 67925855  | 155176058 |                      |         |                     |                     |             |              | 49        |
| **REGION                   | 50 | 1                      | 290379922   | 290379922 | 290379922 |                      |         |                     |                     |             |              | 50        |
| **TOTALS                   |    | 173                    | 14238540145 | 82303700  |           |                      |         |                     |                     |             |              |           |
| DL/I CALLS                 |    |                        |             |           |           |                      |         |                     |                     | IWT/CALL(C) |              |           |
| **REGION                   | 1  | 60                     | 264626241   | 4410437   | 88981490  | 263813671            | 4396894 | 88970053            | 0.76                | 247A,B,C    |              |           |
| **REGION                   | 2  | 223                    | 230505269   | 1033655   | 61048758  | 227368742            | 1019590 | 61011153            | 0.73                | 237A,B,C    |              |           |
| **REGION                   | 3  | 29                     | 257704383   | 8886358   | 69000514  | 257704383            | 8886358 | 69000514            | 0.00                | 98A,B,C     |              |           |
| **REGION                   | 4  | 792                    | 313735347   | 396130    | 52439653  | 310609035            | 392183  | 52439653            | 0.22                | 180A,B,C    |              |           |
| **REGION                   | 49 | 592                    | 262886317   | 444064    | 30202068  | 234394017            | 395935  | 30159782            | 2.46                | 177A,B,C    |              |           |
| **REGION                   | 50 | 36                     | 242591451   | 6738651   | 48651260  | 242591451            | 6738651 | 48651260            | 0.00                | 289A,B,C    |              |           |
| **TOTALS                   |    | 18632                  | 12386905286 | 664818    |           | 12024562411          | 645371  |                     | 0.97                |             |              |           |
| IDLE FOR INTENT            |    |                        |             |           |           |                      |         |                     |                     |             |              |           |
| CHECKPOINT                 |    |                        |             |           |           |                      |         |                     |                     |             |              |           |
| REGION OCCUPANCY           |    |                        |             |           |           |                      |         |                     |                     |             |              |           |
| **REGION                   | 1  | 100.0%                 |             |           |           |                      |         |                     |                     |             |              |           |

```

**REGION 2 100.0%
**REGION 3 100.0%
**REGION 4 100.0%
**REGION 48 100.0%
**REGION 49 100.0%
**REGION 50 100.0%

```

The following example shows a sample of the region-wait report.

```

IMS MONITOR *****REGION IWAIT**** TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0023
*****WAIT TIME*****
**REGION 5 OCCURRENCES TOTAL MEAN MAXIMUM FUNCTION MODULE DISTRIBUTION
NUMBER
SCHEDULING + TERMINATION
SUB-TOTAL
TOTAL
DL/I CALLS
11 181816 16528 24375 DD=IMMSTR2A DBH 117
8 112831 14103 17846 DD=IMMSTR1A DBH 118
5 85460 17092 33717 DD=IMMSTR3A DBH 119
5 58420 11684 14643 DD=IMINDEXA VBH 120
1 4160 4160 4160 INT=DDLTRN24 BLR-64BIT 35
1 3623 3623 3623 PSB=BMPFPE02 BLR-64BIT 36
12 173866 14488 22152 DD=PRODCNTA VBH 121
3 100576 33525 68373 DD=IMMSTR2B DBH 428
1 17921 17921 17921 DD=IMMSTR3B DBH 429
1 17195 17195 17195 DD=IMMSTR1B DBH 430
1 13577 13577 13577 DD=IMINDEXB VBH 431
3 49928 16642 20396 DD=PRODCNTB VBH 432
4 10973 2743 2787 DD=ITEMACTB DBH 453
2 37680 18840 27664 DD=IAINDEXB VBH 454
49 1500067 30613 138284 DD=INVENTRA DBH 472
23 345595 15025 27613 DD=VENDORDA VBH 473
1 342952 342952 342952 PI=VENDORDA...1 498
1 14612 14612 14612 PI=VNSINDXA...1 499
6 69203 11533 19492 DD=VNSINDXA VBH 500
TOTAL
136 3132672 23034

```

The following example shows a sample of the programs-by-region report.

```

IMS MONITOR *****PROGRAMS BY REGION***** TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0069
(A) (B)
ELAPSED EXECUTION TIME SCHEDULING END TO FIRST CALL DISTRIBUTION
OCCURRENCES TOTAL MEAN MAXIMUM TOTAL MEAN MAXIMUM NUMBER
**REGION 1
1
PROGSC6D 1 290146255 290146255 290146255 15479797 15479797 15479797 885A,B
REGION TOTALS 1 290146255 290146255 15479797 15479797
**REGION 2
2
PROGIT8C 1 252290108 252290108 252290108 22376350 22376350 22376350 889A,B
REGION TOTALS 1 252290108 252290108 22376350 22376350
**REGION 3
3
PROGTS1C 1 259496970 259496970 259496970 15169488 15169488 15169488 893A,B
REGION TOTALS 1 259496970 259496970 15169488 15169488
**REGION 4
4
PROGPS3D 1 322812716 322812716 322812716 48146258 48146258 48146258 897A,B
REGION TOTALS 1 322812716 322812716 48146258 48146258
**REGION 5
5
PROGPS3A 2 62893103 31446551 40693590 5435 2717 2862 901A,B
PROGTS1B 1 61794787 61794787 61794787 2790 2790 2790 1271A,B
PROGPS3B 1 18294458 18294458 18294458 3104 3104 3104 1350A,B
PROGIT2B 1 36095342 36095342 36095342 2731 2731 2731 1363A,B
PROGSC2A 1 93902771 93902771 93902771 1667791 1667791 1667791 1401A,B
REGION TOTALS 6 272980461 45496743 1681851 280308
**REGION 6
6
PROGIT1B 2 39000315 19500157 23703429 5286 2643 2801 905A,B
PROGTS1B 1 34293636 34293636 34293636 3136 3136 3136 1207A,B
PROGPS3A 1 51887767 51887767 51887767 2534 2534 2534 1278A,B
PROGPS3B 2 67375031 33687515 40291430 17210570 8605285 17213287 1328A,B
PROGIT8A 1 69132416 69132416 69132416 3291 3291 3291 1359A,B
PROGSC4A 1 30165017 30165017 30165017 2571 2571 2571 1433A,B
REGION TOTALS 8 291854182 36481772 17193752 2149219
**REGION 7
7
PROGSC2B 1 269618583 269618583 269618583 5047875 5047875 5047875 909A,B
REGION TOTALS 1 269618583 269618583 5047875 5047875
**REGION 8
8

```

|               |   |           |          |          |         |         |         |         |
|---------------|---|-----------|----------|----------|---------|---------|---------|---------|
|               | 8 |           |          |          |         |         |         |         |
| PROGIT8A      | 1 | 5181039   | 5181039  | 5181039  | 2928    | 2928    | 2928    | 913A,B  |
| PROGPS3A      | 1 | 27304257  | 27304257 | 27304257 | 3350    | 3350    | 3350    | 1132A,B |
| PROGSC4B      | 1 | 37286872  | 37286872 | 37286872 | 3009    | 3009    | 3009    | 1255A,B |
| PROGIT2A      | 1 | 36902995  | 36902995 | 36902995 | 2850    | 2850    | 2850    | 1298A,B |
| PROGIT1B      | 1 | 30407479  | 30407479 | 30407479 | 2565    | 2565    | 2565    | 1336A,B |
| PROGIT1A      | 3 | 109875360 | 36625120 | 45190114 | 4279008 | 1426336 | 4272096 | 1357A,B |
| PROGIT8B      | 1 | 23405220  | 23405220 | 23405220 | 2679    | 2679    | 2679    | 1395A,B |
| REGION TOTALS | 9 | 270363222 | 30040358 |          | 4296389 | 477376  |         |         |

## Detecting database processing intent conflicts

The IMS Monitor records the intervals when a region is in an idle state waiting to update a database owned exclusively by another already scheduled application program.

You can see the total, maximum, and average idle times in IDLE FOR INTENT following the DL/I calls. The elapsed time during the unsuccessful scheduling of a program in that region is included in the summary line times for that region.

The region can fail to be scheduled even when ownership of that database is released. The number of times processing is held up by intent failure is separately tallied under the title INTENT FAILURE SUMMARY. The report is illustrated in the following example. In this report you can see which PSBs are in conflict because of exclusive intent for a segment type and the database name in question.

| INTENT FAILURE SUMMARY |          |             |
|------------------------|----------|-------------|
| PSBNAME                | DMBNAME  | OCCURRENCES |
| SSTPSBNM               | SSTDMBNM | 1           |
| TOTAL                  |          | 1           |

## Examining the effects of checkpoints

The checkpoint line of the Region Summary report at the end of the region-by-region summary, shows: the number of times that a system checkpoint was taken during the monitor interval, the elapsed times, and the not-wait times.

Checkpoint processing can be initiated by the control program at a specified frequency determined by the number of records placed on the system log. Other checkpoints can be caused by operator commands.

The wait time experienced during checkpoints is reported at the end of the first region summary on the Region Wait report. You can detect delays for each combination of DD name and module code. Typical entries here are for the message queue data sets and the restart data set. If a wait for storage is the cause, the entry under the FUNCTION column is STG.= followed by the identification of the pool.

## Measuring region occupancy

A measure of region activity is the percentage of region occupancy. This is broadly the ratio of the elapsed time a region is performing processing to the trace interval.

The region occupancy time does not include those times when no messages are available, when the block loading is delayed, or when the PSB cannot be used. The last section in the Region Summary report lists all active regions for which timed events were collected and shows the calculated percentages of region occupancy.

## Monitoring application program elapsed time

The IMS Monitor can record measurements of elapsed times for each transaction and scheduling of an application program. It does this during the monitored interval while other programs are executing concurrently.

Elapsed times are calculated from the start of the first DL/I (or other) call to the end of that program. You can distinguish between time spent in application code and in DL/I processing. The following figure illustrates the event intervals.




Figure 88. Event intervals for time in application code and DL/I processing

Within the elapsed time for a DL/I call, the wait time to obtain segment data is recorded separately. Similarly, within the elapsed time for an external subsystem call, the processing time in the external subsystem is recorded separately as the wait time. For regions designated as wait-for-input, the time spent waiting for input messages is excluded from values shown in the reports. The application processing (A/P) time includes many kinds of subsidiary service beyond the machine cycles expended by the program object code—such as subroutine loading, I/O to z/OS data sets, and any overlay processing. If the program is waiting to be dispatched or requires paging before it can use real storage, these delays are also accounted for in application program processing time. Because a program can execute many transactions for each schedule, the elapsed time from schedule to first call is recorded separately. This time covers the initialization performed by the application program and includes the time for loading the program.

The elapsed times are given in the Program Summary report. The following example is a sample of the report. Programs are identified by their PSB name on individual lines in the report. Each line gives a summary of the activity for that PSB during the measured interval. The total number of schedules, DL/I calls, transactions completed (dequeued), and waits for DL/I call I/O and external subsystem processing are given. The report line gives calculated average times for:



- Elapsed time per schedule
- Processor time per schedule
- Schedule to first DL/I call per schedule
- Elapsed time per transaction

The report also includes:

- Frequencies for calls per transaction
- I/O waits per DL/I call
- Waits per external subsystem call
- Transactions dequeued per schedule

A TOTALS line summarizes all activity for the PSBs active during the monitored interval. (The PSB DUMMY line reconciles any incomplete scheduling caused by a region stopping during scheduling or for a program that experiences a pseudoabend.)

| IMS MONITOR |  | ****PROGRAM SUMMARY**** |  |  |  | TRACE START 1993 130 5:55:15 |  |                  |  | TRACE STOP 1993 130 5:59:49 |  |           |  | PAGE 0075 |  |
|-------------|--|-------------------------|--|--|--|------------------------------|--|------------------|--|-----------------------------|--|-----------|--|-----------|--|
|             |  |                         |  |  |  |                              |  | (A).....(B)..... |  | (A).....(B).....            |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  | CPU              |  | ELAPSED                     |  | SCHED. TO |  | ELAPSED   |  |
|             |  |                         |  |  |  |                              |  | TIME             |  | TIME                        |  | 1ST CALL  |  | TIME      |  |
|             |  |                         |  |  |  |                              |  | /SCHED.          |  | /SCHED.                     |  | /SCHED.   |  | /TRANS.   |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |
|             |  |                         |  |  |  |                              |  |                  |  |                             |  |           |  |           |  |

You can use the Call-Summary report to examine the detail of the call processing for each program. The report is itemized by type of call and summarized for the monitor interval. An extract from the multipage output is given in the following example. The calls using an I/O PCB are given first and sub-totaled. Then, the total calls of each type, against each database PCB and each external subsystem, are listed. The PSB TOTAL line marks the end of data for each program.

| IMS MONITOR |     |          |      |           |         | *****CALL SUMMARY***** |           |       |        | TRACE START 1993 130 5:55:15 |                    |                 |                     | TRACE STOP 1993 130 5:59:49 |                 |  |  | PAGE 0186 |  |
|-------------|-----|----------|------|-----------|---------|------------------------|-----------|-------|--------|------------------------------|--------------------|-----------------|---------------------|-----------------------------|-----------------|--|--|-----------|--|
| PSB NAME    |     | PCB NAME |      | CALL FUNC | LEV NO. | SEGMENT                | STAT CODE | CALLS | IWAITS | (C) IWAITS/ CALL             | (A) ..ELAPSED MEAN | TIME... MAXIMUM | (B) .NOT IWAIT MEAN | TIME.. MAXIMUM              | DISTRIB. NUMBER |  |  |           |  |
| PROGSC6B    | I/O | PCB      | ISRT | ( )       |         |                        |           | 138   | 0      | 0.00                         | 372                | 1240            | 372                 | 1240                        | 598A,B,C        |  |  |           |  |
|             |     |          | GU   | ( )       |         |                        |           | 134   | 133    | 0.99                         | 2600917            | 20974615        | 2587532             | 20962866                    | 602A,B,C        |  |  |           |  |
|             |     |          | (GU) | ( )       |         |                        |           | 3     | 0      | 0.00                         | 15                 | 16              | 15                  | 16                          | 716A,B,C        |  |  |           |  |



|  |          |                   |      |      |      |          |          |          |          |          |
|--|----------|-------------------|------|------|------|----------|----------|----------|----------|----------|
|  |          | ASRT ( )          | 3    | 0    | 0.00 | 330      | 333      | 330      | 333      | 869A,B,C |
|  |          | GU ( )            | 2    | 1    | 0.50 | 17639806 | 21219588 | 17634776 | 21209529 | 870A,B,C |
|  |          | I/O PCB SUBTOTAL  |      |      |      |          |          |          |          |          |
|  |          |                   | 280  | 134  | 0.47 | 1370910  |          | 1364469  |          |          |
|  | INVENTRB | DLET (03)IN060SUP | 138  | 0    | 0.00 | 813      | 1289     | 813      | 1289     | 599A,B,C |
|  |          | GNP (03)IN060SUP  | 138  | 7    | 0.05 | 2112     | 112589   | 1047     | 112589   | 600A,B,C |
|  |          | GU (01)IN010PAR   | 138  | 254  | 1.84 | 29511    | 75356    | 1195     | 19229    | 601A,B,C |
|  |          | DL/I PCB SUBTOTAL |      |      |      |          |          |          |          |          |
|  |          |                   | 414  | 261  | 0.63 | 10812    |          | 1018     |          |          |
|  |          | PSB TOTAL         |      |      |      |          |          |          |          |          |
|  |          |                   | 694  | 395  | 0.56 | 559555   |          | 551114   |          |          |
|  | PROGSC2A | I/O PCB           |      |      |      |          |          |          |          |          |
|  |          | ISRT ( )          | 118  | 0    | 0.00 | 381      | 1496     | 381      | 1496     | 603A,B,C |
|  |          | GU ( )            | 114  | 284  | 2.49 | 3304809  | 21784513 | 3164423  | 21664181 | 632A,B,C |
|  |          | (GU) ( )          | 2    | 0    | 0.00 | 17       | 18       | 17       | 18       | 781A,B,C |
|  |          | ASRT ( )          | 3    | 0    | 0.00 | 367      | 444      | 367      | 444      | 871A,B,C |
|  |          | GU ( )            | 2    | 5    | 2.50 | 19931897 | 20045206 | 19799530 | 19925277 | 872A,B,C |
|  |          | I/O PCB SUBTOTAL  |      |      |      |          |          |          |          |          |
|  |          |                   | 239  | 289  | 1.20 | 1743339  |          | 1675270  |          |          |
|  | LOGVENDA | REPL (03)IN040SLQ | 118  | 0    | 0.00 | 268      | 804      | 268      | 804      | 604A,B,C |
|  |          | GNP (03)IN040SLQ  | 118  | 5    | 0.04 | 899      | 16995    | 218      | 305      | 605A,B,C |
|  |          | REPL (02)VN030PAR | 826  | 0    | 0.00 | 805      | 1578     | 805      | 1578     | 606A,B,C |
|  |          | GNP (02)VN030PAR  | 826  | 873  | 1.05 | 19321    | 94521    | 456      | 1363     | 607A,B,C |
|  |          | REPL (01)VN020REO | 118  | 58   | 0.49 | 8879     | 48076    | 832      | 1682     | 623A,B,C |
|  |          | GU (01)VN020REO   | 118  | 195  | 1.65 | 31688    | 360775   | 1300     | 1746     | 625A,B,C |
|  |          | DL/I PCB SUBTOTAL |      |      |      |          |          |          |          |          |
|  |          |                   | 2124 | 1131 | 0.53 | 10145    |          | 636      |          |          |
|  |          | PSB TOTAL         |      |      |      |          |          |          |          |          |
|  |          |                   | 2363 | 1420 | 0.60 | 185445   |          | 170013   |          |          |
|  | PROGSC2D | I/O PCB           |      |      |      |          |          |          |          |          |
|  |          | ISRT ( )          | 14   | 0    | 0.00 | 377      | 621      | 377      | 621      | 608A,B,C |
|  |          | GU ( )            | 14   | 36   | 2.57 | 22360408 | 52048566 | 22221852 | 51901313 | 634A,B,C |
|  |          | I/O PCB SUBTOTAL  |      |      |      |          |          |          |          |          |
|  |          |                   | 28   | 36   | 1.28 | 11180393 |          | 11111115 |          |          |
|  | LOGVENDD | REPL (03)IN040SLQ | 14   | 0    | 0.00 | 263      | 328      | 263      | 328      | 609A,B,C |
|  |          | GNP (03)IN040SLQ  | 14   | 1    | 0.07 | 1407     | 16889    | 223      | 307      | 610A,B,C |
|  |          | REPL (02)VN030PAR | 98   | 0    | 0.00 | 820      | 1015     | 820      | 1015     | 611A,B,C |
|  |          | (GU) (..).....    | 1    | 0    | 0.00 |          |          |          |          |          |
|  | APOL1    | I/O PCB           |      |      |      |          |          |          |          |          |
|  |          | I/O PCB SUBTOTAL  |      |      |      |          |          |          |          |          |
|  |          |                   | 1    | 0    | 0.00 |          |          |          |          |          |
|  | OTMDEST1 | ICAL (..).....    | 1    | 0    | 0.00 |          |          |          |          |          |
|  |          | DL/I PCB SUBTOTAL |      |      |      |          |          |          |          |          |
|  |          |                   | 1    | 0    | 0.00 |          |          |          |          |          |
|  |          | PSB TOTAL         |      |      |      |          |          |          |          |          |
|  |          |                   | 2    | 0    | 0.00 |          |          |          |          |          |

## Monitoring I/O for application program DL/I Calls

The IMS Monitor report shows the total number of I/O occurrences and the total time the occurrences took for each application program executed during a monitored interval.

The Program-I/O report gives these two totals for all PSBs active during the monitored interval and includes the detailed breakdown of the I/O wait time as it was incurred by each PCB used by the program.

The report shows any contention experienced during application program processing. Each type of conflict and the number of times it occurred are recorded for each I/O PCB or database PCB. The report shows the total wait time, the highest wait experienced, and the average time. Subtotals are given for each PCB under a PSB, and for all PCBs under each PSB.

The DDN/FUNC column lists the data set DD name. The MODULE column uses a code to indicate the source of the contention. The types of conflicts and codes are shown as follows.

- **Message handling**

**Code**    **Conflict**

**DBH**    OSAM I/O for message queues

**MFS**    MFS format library directory

**PMM**    Message format buffer pool space or control block I/O

**QMG**    Message queue management

- **Scheduling**

**Code    Conflict**

**BLR**    Load/read from ACBLIB

**MSC**    MPP region initialization

**SMN**    Virtual storage management

- **Database access**

**Code    Conflict**

**DBH**    OSAM I/O

**DLE**    DL/I functions

**VBH**    VSAM interface

**(Physical segment code)**

Program isolation

For external subsystem calls, the elapsed time to complete the processing is considered wait time. The DDN/FUNC column indicates the external subsystem call function, as shown as follows:

- **External subsystems**

**Code    Subsystem call function**

**AB0**    ABORT

**CT0**    Create thread

**D50**    Terminate identify or thread, signoff

**D80**    INIT

**I30**    Identify, command, echo, terminate

**I30**    Identify, terminate subsystem

**I50**    INIT

**I60**    Resolve-in-doubt

**PR0**    Subsystem-not-operational

**P10**    Commit prepare (Phase 1)

**P20**    Commit continue (Phase 2)

**SO0**    Signon

**SI0**    Identify

The following example show a sample of the report.

| IMS MONITOR |          | ****PROGRAM I/O****  |          | TRACE START 1993 022 |         | 14:00:18 |        | TRACE STOP |  | 1993 022 |  | 14:02:20 |  | PAGE 0088 |  |
|-------------|----------|----------------------|----------|----------------------|---------|----------|--------|------------|--|----------|--|----------|--|-----------|--|
|             |          | .....IWAIT TIME..... |          |                      |         |          |        |            |  |          |  |          |  |           |  |
| PSBNAME     | PCB NAME | IWAITS               | TOTAL    | MEAN                 | MAXIMUM | DDN/FUNC | MODULE |            |  |          |  |          |  |           |  |
| PROGHR1A    | I/O PCB  | 122                  | 2341116  | 19189                | 70795   | HOTELDBA | DBH    |            |  |          |  |          |  |           |  |
|             |          | 34                   | 24177936 | 711115               | 3950160 | **W F I  |        |            |  |          |  |          |  |           |  |
|             |          | 40                   | 23652665 | 591316               | 2668917 | **W F I  |        |            |  |          |  |          |  |           |  |
|             |          | 5                    | 67613    | 13522                | 21214   | SHMSG    | QMG    |            |  |          |  |          |  |           |  |
|             |          | 4                    | 110363   | 27590                | 60486   | QBLKS    | QMG    |            |  |          |  |          |  |           |  |
| PCB TOTAL   |          | 131                  | 2519092  | 19229                |         |          |        |            |  |          |  |          |  |           |  |
| PSB TOTAL   |          | 305                  | 6725063  | 20049                |         |          |        |            |  |          |  |          |  |           |  |

|                                 |        |          |                      |          |                     |                    |
|---------------------------------|--------|----------|----------------------|----------|---------------------|--------------------|
| PROGDE1A TRMNALDA               | 20     | 624677   | 31233                | 68252    | TRMNALDA            | VBH                |
|                                 | 1      | 275811   | 275811               | 275811   | PI TRMNALDA....     |                    |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
| I/O PCB                         | 21     | 900488   | 42880                |          |                     |                    |
|                                 | 16     | 488812   | 30550                | 79980    | TRMNALDA            | VBH                |
|                                 | 1      | 16118    | 16118                | 16118    | SHMSG               | QMG                |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
| TABLEDBA                        | 17     | 504930   | 29701                |          |                     |                    |
|                                 | 16     | 290471   | 18154                | 33254    | TABLEDA             | DBH                |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 16     | 290471   | 18154                |          |                     |                    |
| PSB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 54     | 1695889  | 31405                |          |                     |                    |
| PROGHR2B HOTELDBB               | 8      | 698384   | 87298                | 184475   | HOTELDBB            | DBH                |
|                                 | 4      | 5820650  | 1455162              | 1455278  | PI HOSINDXB....     |                    |
|                                 | 4      | 4481024  | 1120256              | 1209075  | PI HOTELDBB....     |                    |
|                                 | 2      | 260817   | 130408               | 232750   | HOSINDOB            | VBH                |
|                                 | 7      | 106623   | 15231                | 16410    | HOSINDXB            | VBH                |
|                                 | 1      | 15366    | 15366                | 15366    | HOTELDBD            | DBH                |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 26     | 11382864 | 437802               |          |                     |                    |
| PSB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 26     | 11382864 | 437802               |          |                     |                    |
| PROGHR2A HOTELDBA               | 17     | 655801   | 38576                | 366108   | HOSINDXA            | VBH                |
|                                 | 73     | 1836721  | 25160                | 82141    | HOTELDBA            | DBH                |
|                                 | 2      | 54663    | 27331                | 41975    | HOTELDBD            | DBH                |
|                                 | 1      | 9887     | 9887                 | 9887     | HOTELDBC            | DBH                |
|                                 | 2      | 851042   | 845635               | 845635   | HOSINDOA            | VBH                |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
| I/O PCB                         | 95     | 3408114  | 35874                |          |                     |                    |
|                                 | 20     | 575847   | 28792                | 74227    | HOTELDBA            | DBH                |
|                                 | 21     | 370390   | 17637                | 43153    | HOSINDXA            | VBH                |
| IMS MONITOR ****PROGRAM I/O**** |        |          |                      |          |                     |                    |
|                                 |        |          | TRACE START 1993 022 | 14:00:18 | TRACE STOP 1993 022 | 14:02:20 PAGE 0089 |
|                                 |        |          | .....IWAIT TIME..... |          |                     |                    |
| PSBNAME PCB NAME                | IWAITS | TOTAL    | MEAN                 | MAXIMUM  | DDN/FUNC            | MODULE             |
| PROGHR2A I/O PCB                | 5      | 4654544  | 930908               | 2020043  | **W F I             |                    |
|                                 | 8      | 32796604 | 4099575              | 9328891  | **W F I             |                    |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 41     | 946237   | 23078                |          |                     |                    |
| PSB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 136    | 4354351  | 32017                |          |                     |                    |
| PROGPS2A LOGIMA                 | 89     | 2046670  | 22996                | 73593    | IMMSTR3A            | VBH                |
|                                 | 612    | 53886417 | 88049                | 185674   | IMMSTR1A            | VBH                |
|                                 | 3      | 44906    | 14968                | 20788    | IMINDEXA            | VBH                |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 704    | 55977993 | 79514                |          |                     |                    |
|                                 | 469    | 11742900 | 25038                | 170337   | COMPOSDA            | DBH                |
|                                 | 329    | 8198418  | 24919                | 91422    | CPINDEXA            | VBH                |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
| I/O PCB                         | 798    | 19941318 | 24989                |          |                     |                    |
|                                 | 3      | 47511    | 15837                | 20806    | SHMSG               | QMG                |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 3      | 47511    | 15837                |          |                     |                    |
| PSB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 1505   | 75966822 | 50476                |          |                     |                    |
| PROGSC6C I/O PCB                | 52     | 2698602  | 51896                | 473763   | INVENTRC            | VBH                |
|                                 | 4      | 70921    | 17730                | 34241    | SHMSG               | QMG                |
|                                 | 3      | 50699    | 16899                | 24724    | QBLKS               | QMG                |
| PCB TOTAL                       |        |          |                      |          |                     |                    |
|                                 | 59     | 2820222  | 47800                |          |                     |                    |
|                                 | 55     | 2666884  | 48488                | 210752   | INVENTRC            | VBH                |

|                                                                                                  |          |        |                      |         |         |                 |        |
|--------------------------------------------------------------------------------------------------|----------|--------|----------------------|---------|---------|-----------------|--------|
|                                                                                                  |          | 50     | 797587               | 15951   | 41706   | ININDEXC        | VBH    |
|                                                                                                  |          | 1      | 119253               | 119253  | 119253  | PI INVENTRC...  | 1      |
|                                                                                                  |          | 1      | 8634                 | 8634    | 8634    | INVENTRB        | VBH    |
|                                                                                                  |          | 2      | 83947                | 41973   | 53936   | INVENTRA        | VBH    |
| PCB TOTAL                                                                                        |          |        |                      |         |         |                 |        |
| PSB TOTAL                                                                                        |          | 109    | 3676305              | 33727   |         |                 |        |
|                                                                                                  |          | 168    | 6496527              | 38669   |         |                 |        |
| PROGHR2D I/O PCB                                                                                 |          |        |                      |         |         |                 |        |
|                                                                                                  |          | 21     | 2285296              | 108823  | 199223  | HOTELDBD        | DBH    |
|                                                                                                  |          | 28     | 762370               | 27227   | 111860  | HOSINDXD        | VBH    |
|                                                                                                  |          | 1      | 11685                | 11685   | 11685   | SHMSG           | QMG    |
| PCB TOTAL                                                                                        |          |        |                      |         |         |                 |        |
| HOTELDBD                                                                                         |          | 50     | 3059351              | 61187   |         |                 |        |
|                                                                                                  |          | 96     | 6279107              | 65407   | 139032  | HOTELDBD        | DBH    |
| MONITOR ****PROGRAM I/O**** TRACE START 1993 022 14:00:18 TRACE STOP 1993 022 14:02:20 PAGE 0090 |          |        |                      |         |         |                 |        |
| PSBNAME                                                                                          | PCB NAME | IWAITS | .....IWAIT TIME..... |         |         |                 |        |
|                                                                                                  |          |        | TOTAL                | MEAN    | MAXIMUM | DDN/FUNC        | MODULE |
| PROGHR2D                                                                                         | HOTELDBD | 31     | 2130585              | 68728   | 769130  | HOSINDXD        | VBH    |
|                                                                                                  |          | 3      | 115999               | 38666   | 56394   | HOTELDBA        | DBH    |
|                                                                                                  |          | 2      | 69833                | 34916   | 43470   | HOTELDBC        | DBH    |
|                                                                                                  |          | 2      | 41430                | 20715   | 28020   | HOSINDOD        | VBH    |
|                                                                                                  |          | 4      | 5515374              | 1378843 | 1458884 | PI HOSINDXD.... |        |
|                                                                                                  |          | 4      | 3997017              | 999254  | 1026228 | PI HOTELDBD.... |        |
| PCB TOTAL                                                                                        |          |        |                      |         |         |                 |        |
|                                                                                                  |          | 142    | 18149345             | 127812  |         |                 |        |
| PSB TOTAL                                                                                        |          | 192    | 21208696             | 110461  |         |                 |        |

The I/O waits for the calls to the I/O PCB, are grouped as the first entries for a PSB. For DL/I calls, the data set for which the I/O took place is indicated under the DDN/FUNC heading, and the module code indicates what type of conflict caused the wait. For external subsystem calls, the function is indicated under the DDN/FUNC heading and the module code indicates the source of the call entry.

Names other than LGMSG and SHMSG can appear in the DDN/FUNC column for I/O PCBs. An example is a checkpoint call issued by an application program (using an I/O PCB) which causes a database buffer to be written.

If the program is designated as wait-for-input and has to wait for the input of the next message, the wait entry is marked \*\*WFI under the DDN/FUNC heading and no entry appears in the MODULE column. The time spent waiting for the next input message is shown under wait time. \*\*WFI entries are shown for information only and their values are not used to compute statistics.

Contention for the same physical segment in a database causes a wait on behalf of program isolation. This is shown in the DDN/FUNC column, on the PCB line, by the entry PIdmb, where dmb is the DMB of the physical data set. The MODULE column identifies the segment type using the physical segment code assigned by DBD generation.

When an application is accessing a database using VSAM as the access method, DL/I calls do not generally result in an I/O wait. A MODULE column entry of VBH indicates that interface to VSAM occurred and there was an I/O wait.

A seemingly unrelated entry can occur under the DDN/FUNC column for a database PCB. An example is a retrieval call to a database (DB-A) that causes a buffer to be purged in order to make room for that retrieved data. If the buffer contents included data belonging to another database (DB-B), the I/O entry in the report shows the DD name for DB-B as being in conflict for PCB access to DB-A.

---

## Monitoring MFS activity

You can obtain a summary of all activity that occurs for management of message format buffer pool use from the Message Format Buffer Pool report.

The report is illustrated in the following example. The data shows the counts at the start and end of the trace interval and their difference.

When message formatting occurs, the appropriate message blocks must reside in the message format buffer pool, a DIF/MID pair for input or a DOF/MOD pair for output. If the blocks are not already in the buffer, I/O to the active IMS.FORMATA/B library must occur. Block retrieval can involve a prior directory lookup, or be direct, using an index kept in the pool.

Many of the counts reveal details of internal event management. When there is no directory entry for a block this implies extra directory lookup I/O.

```
***I M S   M O N I T O R***  BUFFER POOL STATISTICS  TRACE START 1993 130   5:55:15  TRACE STOP  1993 130   5:59:49  PAGE 0007

      M E S S A G E   F O R M A T   B U F F E R   P O O L

                                     5:55:15      5:59:49
                                     START TRACE    END   TRACE      DIFFERENCE

NUMBER OF P/F REQUESTS                0              0              0
NUMBER OF I/F REQUESTS               18             20              2
NUMBER OF I/F I/O'S                   2              2              0
NUMBER OF TIMES POOL COMPRESS WOULD BE SUCCESSFUL  0              0              0
NUMBER OF DIRECTORY I/O OPERATIONS    2              2              0
NUMBER OF TIMES BLOCK WASHED FOR FRE   0              0              0
NUMBER OF TIMES P/F REQUEST IGNORED    0              0              0
NUMBER OF F/B REQUESTS               18             20              2
NUMBER OF TIMES F/B REQUEST IGNORED    0              0              0
NUMBER OF TIMES I/F ON F/B QUEUE      16             18              2
NUMBER OF TIMES I/F ON I/F QUEUE       0              0              0
NUMBER OF TIMES F/B ON I/F QUEUE      18             20              2
NUMBER OF TIMES P/F ON I/F QUEUE       0              0              0
NUMBER OF TIMES P/F ON F/B QUEUE       0              0              0
NUMBER OF TIMES THERE WAS NO DIR ENTR FOR A BLOCK  0              0              0
NUMBER OF TIMES I/O ERRORS POINT OR READ MACRO    0              0              0
NUMBER OF IMMEDIATE I/O REQUESTS WAITED DUE TO MAXIMUM I/O  0              0              0
NUMBER OF REQUESTS SATISFIED BY INDEX/DYNAMIC DIRECTORY  0              0              0

QUOTIENT :  IMMEDIATE FETCH I/O'S + DIRECTORY I/O'S OPERATIONS =    0.00

      _____
                        TOTAL NUMBER OF TRANSACTIONS
```

---

## Monitoring message queue handling

A key resource that directly affects the efficiency of transaction processing is the message queue pool and the management of the I/O to the message queues. You can examine the activity by looking at the Message Queue Pool report.

The following example illustrates the message-queue-pool report. Counts of activities are given at start and end of the trace interval and as the differences between start and end numbers.

```
***I M S   M O N I T O R***  BUFFER POOL STATISTICS  TRACE START 1993 130   5:55:15  TRACE STOP  1993 130   5:59:49  PAGE 0002

      M E S S A G E   Q U E U E   P O O L

                                     5:55:15      5:59:49
                                     START TRACE    END   TRACE      DIFFERENCE

NUMBER OF LOCATE CALLS FROM QMGR       54204           68436          14232
NUMBER OF RECORD RELEASE CALLS FROM QMGR 16431           20738           4307
NUMBER OF LOCATE AND ALTER CALLS FROM QMGR 131593          164744          33151
NUMBER OF REQUESTS TO PURGE THE Q POOL      2              2              0
NUMBER OF ADDRESS TO DRRN TRANSLATION REQUESTS 21351          27076           5725
```

|                                                     |       |       |       |
|-----------------------------------------------------|-------|-------|-------|
| NUMBER OF REQUESTS TO WAIT FROM QMGR                | 0     | 0     | 0     |
| NUMBER OF READ REQUESTS                             | 962   | 962   | 0     |
| NUMBER OF WRITE REQUESTS(TOTAL)                     | 499   | 499   | 0     |
| NUMBER OF WRITES DONE BY PURGE                      | 499   | 499   | 0     |
| NUMBER OF WAITS FOR PURGE COMPLETION                | 1     | 1     | 0     |
| NUMBER OF WAITS BECAUSE NO BUFFER AVAILABLE         | 0     | 0     | 0     |
| NUMBER OF WAITS FOR OTHER DECB TO READ THIS BUFFER  | 823   | 823   | 0     |
| NUMBER OF WAITS FOR OTHER DECB TO WRITE THIS BUFFER | 0     | 0     | 0     |
| NUMBER OF WAITS FOR CONFLICTING END DEQ BUFFER REQ  | 0     | 0     | 0     |
| NUMBER OF PSBS UNCHAINED FROM BUFFERS               | 0     | 0     | 0     |
| NUMBER OF CALLS TO QMGR.(TOTAL)                     | 48164 | 62213 | 14049 |
| NUMBER OF CALLS TO REPOSITION A LOST BUFFER         | 0     | 0     | 0     |
| NUMBER OF CALLS TO ENQ A MESSAGE                    | 10583 | 13441 | 2858  |
| NUMBER OF CALLS TO DEQ ONE OR MORE MESSAGE          | 6321  | 7767  | 1446  |
| NUMBER OF CALLS TO CANCEL INPUT OR OUTPUT           | 119   | 121   | 2     |

QUOTIENT : TOTAL NUMBER OF OSAM READS + OSAM WRITES + ALL IWAITS = 0.00

TOTAL NUMBER OF TRANSACTIONS

## Detecting checkpoint effects

When a checkpoint command specifies SNAPQ, the current status of all message queues is written to the system log. This prevents any message handling on behalf of queue management.

The General Iwait Time Events records the wait time incurred by the SNAPQ. The following example shows the activity on the summary line QMGR SNAPQ CHECK. The number of occurrences is given with the total, average, and maximum wait times.

```

IMS MONITOR ** GENERAL REPORTS **   TRACE START 1993 130   5:55:15   TRACE STOP 1993 130   5:59:49   PAGE 0009
GENERAL IWAIT TIME EVENTS

EVENT          OCCURRENCES  .....IWAIT TIME.....  DISTRIBUTION
IWAITS                                     TOTAL      MEAN      MAXIMUM      NUMBER
-----
QMGR SNAPQ CHECK          0          0          0          0          0
REGION AND JOBNAME REPORT
REG. NO.  JOB NAME
-----
1      MPR1A100
2      MPR1A209
3      MPR1A210
4      MPR1A211
5      MPR1A103
6      MPR1A101
7      MPR1A115
8      MPR1A116
9      MPR1A216
10     MPR1A200
11     MPR1A217
12     MPR1A119
13     MPR1A218
14     MPR1A219
15     MPR1A104
16     MPR1A220
17     MPR1A203
18     MPR1A123
19     MPR1A222
20     MPR1A105
21     MPR1A124
22     MPR1A223
23     MPR1A107
24     MPR1A224
25     MPR1A106
26     MPR1A206
27     MPR1A205
28     MPR1A108
29     MPR1A109
30     MPR1A208
31     MPR1A111
32     MPR1A112
33     MPR1A113
34     MPR1A204
35     MPR1A114

```

36 MPR1A102  
 48 MPR1A121  
 49 MPR1A122  
 50 MPR1A221

## Transaction Queuing report

In addition to monitoring the efficiency of message handling, you can monitor the service provided for each application, by looking at the size of the transaction queues at each scheduling of their processing programs.

The Transaction-Queuing report shown in the following example records, for each transaction, the minimum, average, and maximum counts at scheduling time. The total number of dequeued transactions (or transactions that have been fully processed) during the monitored interval is given for each transaction code. The average number of transactions processed for each scheduling is given in the DEQUEUED MEAN column.

| IMS MONITOR *****TRANSACTION QUEUING***** TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0181 |                    |                   |                       |                                                  |                         |                        |         |
|--------------------------------------------------------------------------------------------------------------|--------------------|-------------------|-----------------------|--------------------------------------------------|-------------------------|------------------------|---------|
| TRANSACTION                                                                                                  | NUMBER<br>DEQUEUED | NUMBER<br>SCHEDS. | ..ON QUEUE<br>MINIMUM | (B)<br>WHEN<br>SCHEDULED.....<br>MEAN<br>MAXIMUM | (A)<br>DEQUEUED<br>MEAN | DISTRIBUTION<br>NUMBER |         |
| SC6X                                                                                                         | 13                 | 1                 | 0                     | 0.00                                             | 0                       | 13.00                  | 883A,B  |
| IT8W                                                                                                         | 17                 | 3                 | 0                     | 0.00                                             | 0                       | 5.66                   | 887A,B  |
| TS1Z                                                                                                         | 16                 | 1                 | 0                     | 0.00                                             | 0                       | 16.00                  | 891A,B  |
| PS3X                                                                                                         | 23                 | 1                 | 0                     | 0.00                                             | 0                       | 23.00                  | 895A,B  |
| PS3Y                                                                                                         | 17                 | 7                 | 0                     | 0.00                                             | 0                       | 2.42                   | 899A,B  |
| IT1V                                                                                                         | 11                 | 6                 | 0                     | 0.00                                             | 0                       | 1.83                   | 903A,B  |
| SC2Z                                                                                                         | 143                | 2                 | 0                     | 0.00                                             | 0                       | 71.50                  | 907A,B  |
| IT8U                                                                                                         | 12                 | 7                 | 0                     | 0.00                                             | 0                       | 1.71                   | 911A,B  |
| PS2W                                                                                                         | 10                 | 1                 | 0                     | 0.00                                             | 0                       | 10.00                  | 915A,B  |
| TS1U                                                                                                         | 12                 | 4                 | 0                     | 0.00                                             | 0                       | 3.00                   | 919A,B  |
| PS3W                                                                                                         | 14                 | 1                 | 0                     | 0.00                                             | 0                       | 14.00                  | 923A,B  |
| IT8Y                                                                                                         | 16                 | 5                 | 0                     | 0.00                                             | 0                       | 3.20                   | 927A,B  |
| IT1W                                                                                                         | 9                  | 1                 | 0                     | 0.00                                             | 0                       | 9.00                   | 929A,B  |
| SC2W                                                                                                         | 9                  | 1                 | 0                     | 0.00                                             | 0                       | 9.00                   | 933A,B  |
| IT2V                                                                                                         | 13                 | 5                 | 0                     | 0.00                                             | 0                       | 2.60                   | 937A,B  |
| IT2W                                                                                                         | 17                 | 6                 | 0                     | 0.00                                             | 0                       | 2.83                   | 941A,B  |
| TS1V                                                                                                         | 9                  | 1                 | 0                     | 0.00                                             | 0                       | 9.00                   | 945A,B  |
| SC2V                                                                                                         | 12                 | 5                 | 0                     | 0.00                                             | 0                       | 2.40                   | 947A,B  |
| TS1W                                                                                                         | 11                 | 1                 | 0                     | 0.00                                             | 0                       | 11.00                  | 949A,B  |
| PS3V                                                                                                         | 13                 | 3                 | 0                     | 0.00                                             | 0                       | 4.33                   | 953A,B  |
| IT1U                                                                                                         | 9                  | 6                 | 0                     | 0.00                                             | 0                       | 1.50                   | 957A,B  |
| SC4U                                                                                                         | 11                 | 5                 | 0                     | 0.00                                             | 0                       | 2.20                   | 962A,B  |
| SC6W                                                                                                         | 10                 | 1                 | 0                     | 0.00                                             | 0                       | 10.00                  | 966A,B  |
| PS2V                                                                                                         | 8                  | 6                 | 0                     | 0.00                                             | 0                       | 1.33                   | 970A,B  |
| IT8X                                                                                                         | 12                 | 1                 | 0                     | 0.00                                             | 0                       | 12.00                  | 974A,B  |
| SC4W                                                                                                         | 10                 | 1                 | 0                     | 0.00                                             | 0                       | 10.00                  | 978A,B  |
| SC6U                                                                                                         | 14                 | 6                 | 0                     | 0.00                                             | 0                       | 2.33                   | 982A,B  |
| IT2Y                                                                                                         | 9                  | 3                 | 0                     | 0.00                                             | 0                       | 3.00                   | 986A,B  |
| SC2X                                                                                                         | 15                 | 1                 | 0                     | 0.00                                             | 0                       | 15.00                  | 990A,B  |
| PS2Y                                                                                                         | 17                 | 2                 | 0                     | 0.00                                             | 0                       | 8.50                   | 994A,B  |
| SC4Y                                                                                                         | 152                | 4                 | 0                     | 0.50                                             | 1                       | 38.00                  | 998A,B  |
| SC2Y                                                                                                         | 106                | 2                 | 0                     | 0.00                                             | 0                       | 53.00                  | 1000A,B |
| IT2X                                                                                                         | 20                 | 1                 | 0                     | 0.00                                             | 0                       | 20.00                  | 1004A,B |
| SC2U                                                                                                         | 15                 | 3                 | 0                     | 0.00                                             | 0                       | 5.00                   | 1008A,B |
| SC4Z                                                                                                         | 123                | 5                 | 0                     | 0.60                                             | 1                       | 24.60                  | 1010A,B |
| TS1X                                                                                                         | 15                 | 1                 | 0                     | 0.00                                             | 0                       | 15.00                  | 1015A,B |
| SC4X                                                                                                         | 19                 | 1                 | 0                     | 0.00                                             | 0                       | 19.00                  | 1019A,B |
| PS2X                                                                                                         | 13                 | 1                 | 0                     | 0.00                                             | 0                       | 13.00                  | 1024A,B |
| PS2Z                                                                                                         | 20                 | 5                 | 0                     | 0.00                                             | 0                       | 4.00                   | 1028A,B |
| SC6Z                                                                                                         | 130                | 1                 | 0                     | 0.00                                             | 0                       | 130.00                 | 1031A,B |
| SC6V                                                                                                         | 10                 | 4                 | 0                     | 0.00                                             | 0                       | 2.50                   | 1035A,B |
| SC6Y                                                                                                         | 143                | 1                 | 0                     | 0.00                                             | 0                       | 143.00                 | 1037A,B |
| IT1X                                                                                                         | 10                 | 1                 | 0                     | 0.00                                             | 0                       | 10.00                  | 1040A,B |
| PS3U                                                                                                         | 19                 | 6                 | 0                     | 0.00                                             | 0                       | 3.16                   | 1131A,B |
| IT2U                                                                                                         | 13                 | 4                 | 0                     | 0.00                                             | 0                       | 3.25                   | 1146A,B |

# Monitoring database buffers

One of the key resources in an online system is the database buffer pool. The efficiency of DL/I call service depends on the presence of the required database logical record in the buffer, so that segment retrieval does not require additional I/O.

This is especially true for HOLD calls with intervening database calls prior to a replace call.

Use the QUERY POOL command to view usage statistics for Fast Path buffer pools.

You can assess the general efficiency of the full-function pool management using the Database Buffer Pool report shown in the following example. The event counts on this report are not specific to a particular database or program but represent the pressure for use of the database pool.

**Related reading:** For more information about the Database Buffer Pool reports, see "Database-Buffer-Pool report" on page 720.

If any of your databases use VSAM as the access method, the IMS Monitor produces a series of reports headed VSAM BUFFER POOL, one for each subpool. The following example shows one of these reports.

| DATA BASE BUFFER POOL                                  |                          |          |        |
|--------------------------------------------------------|--------------------------|----------|--------|
|                                                        | FIX PREFIX/BUFFERS       |          | Y/Y    |
|                                                        | SUBPOOL ID               |          | 004K   |
|                                                        | SUBPOOL BUFFER SIZE      |          | 4096   |
|                                                        | TOTAL BUFFERS IN SUBPOOL |          | 1000   |
|                                                        | 17:08:15                 | 17:10:16 |        |
| NUMBER OF LOCATE-TYPE CALLS                            | 1117674                  | 1676213  | 558539 |
| NUMBER OF REQUESTS TO CREATE NEW BLOCKS                | 0                        | 0        | 0      |
| NUMBER OF BUFFER ALTER CALLS                           | 215874                   | 322936   | 107062 |
| NUMBER OF PURGE CALLS                                  | 25077                    | 37454    | 12377  |
| NUMBER OF LOCATE-TYPE CALLS, DATA ALREADY IN OSAM POOL | 870306                   | 1301187  | 430881 |
| NUMBER OF BUFFERS SEARCHED BY ALL LOCATE-TYPE CALLS    | 1258247                  | 1886843  | 628596 |
| NUMBER OF READ I/O REQUESTS                            | 238165                   | 360260   | 122095 |
| NUMBER OF SINGLE BLOCK WRITES BY BUFFER STEAL ROUTINE  | 0                        | 0        | 0      |
| NUMBER OF BLOCKS WRITTEN BY PURGE                      | 95057                    | 142413   | 47356  |
| NUMBER OF LOCATE CALLS WAITED DUE TO BUSY ID           | 780                      | 1297     | 517    |
| NUMBER OF LOCATE CALLS WAITED DUE TO BUFFER BUSY WRT   | 0                        | 0        | 0      |
| NUMBER OF LOCATE CALLS WAITED DUE TO BUFFER BUSY READ  | 0                        | 0        | 0      |
| NUMBER OF BUFFER STEAL/PURGE WAITED FOR OWNERSHIP RLSE | 178                      | 261      | 83     |
| NUMBER OF BUFFER STEAL REQUESTS WAITED FOR BUFFERS     | 0                        | 0        | 0      |
| TOTAL NUMBER OF I/O ERRORS FOR THIS SUBPOOL            | 0                        | 0        | 0      |
| NUMBER OF BUFFERS LOCKED DUE TO WRITE ERRORS           | 0                        | 0        | 0      |
| QUOTIENT : TOTAL NUMBER OF OSAM READS + OSAM WRITES =  | 6.98                     |          |        |
| TOTAL NUMBER OF TRANSACTIONS                           |                          |          |        |

The following example shows a sample of the VSAM-Buffer-Pool report.

|                                          |                           |  |       |
|------------------------------------------|---------------------------|--|-------|
| ***IMS MONITOR*** BUFFER POOL STATISTICS |                           |  |       |
| VSAM BUFFER POOL                         |                           |  |       |
|                                          | FIX INDEX/BLOCK/DATA      |  | N/Y/N |
|                                          | SHARED RESOURCE POOL ID   |  | VPL1  |
|                                          | SHARED RESOURCE POOL TYPE |  | D     |
|                                          | SUBPOOL ID                |  | 2     |
|                                          | SUBPOOL BUFFER SIZE       |  | 4096  |
|                                          | NUMBER HIPERSPACE BUFFERS |  | 50    |
|                                          | TOTAL BUFFERS IN SUBPOOL  |  | 1000  |



|                                                         | 17:08:15                     | 17:10:16 |       |
|---------------------------------------------------------|------------------------------|----------|-------|
| NUMBER OF RETRIEVE BY RBA CALLS RECEIVED BY BUF HNDLR   | 152                          | 330      | 178   |
| NUMBER OF RETRIEVE BY KEY CALLS                         | 117780                       | 178424   | 60644 |
| NUMBER OF LOGICAL RECORDS INSERTED INTO ESDS            | 132                          | 310      | 178   |
| NUMBER OF LOGICAL RECORDS INSERTED INTO KSDS            | 6460                         | 9853     | 3393  |
| NUMBER OF LOGICAL RECORDS ALTERED IN THIS SUBPOOL       | 0                            | 0        | 0     |
| NUMBER OF TIMES BACKGROUND WRITE FUNCTION INVOKED       | 0                            | 0        | 0     |
| NUMBER OF SYNCHRONIZATION CALLS RECEIVED                | 18566                        | 27923    | 9357  |
| NUMBER OF WRITE ERROR BUFFERS CURRENTLY IN THE SUBPOOL  | 0                            | 0        | 0     |
| LARGEST NUMBER OF WRITE ERRORS IN THE SUBPOOL           | 0                            | 0        | 0     |
| NUMBER OF VSAM GET CALLS ISSUED                         | 124648                       | 189220   | 64572 |
| NUMBER OF VSAM SCHBFR CALLS ISSUED                      | 0                            | 0        | 0     |
| NUMBER OF TIMES CTRL INTERVAL REQUESTED ALREADY IN POOL | 33662                        | 51088    | 17426 |
| NUMBER OF CRTL INTERVALS READ FROM EXTERNAL STORAGE     | 91169                        | 138505   | 47336 |
| NUMBER OF VSAM WRITES INITIATED BY IMS                  | 6022                         | 9251     | 3229  |
| NUMBER OF VSAM WRITES TO MAKE SPACE IN THE POOL         | 0                            | 0        | 0     |
| NUMBER OF VSAM READS FROM HIPERSPACE BUFFERS            | 0                            | 0        | 0     |
| NUMBER OF VSAM WRITES FROM HIPERSPACE BUFFERS           | 50                           | 50       | 0     |
| NUMBER OF FAILED VSAM READS FROM HIPERSPACE BUFFERS     | 0                            | 0        | 0     |
| NUMBER OF FAILED VSAM WRITES FROM HIPERSPACE BUFFERS    | 0                            | 0        | 0     |
| QUOTIENT : TOTAL NUMBER OF VSAM READS + VSAM WRITES =   | 2.08                         |          |       |
|                                                         | TOTAL NUMBER OF TRANSACTIONS |          |       |

## Monitoring line activity

You can obtain a summary of all occurrences of activity for each node that handles message traffic during the monitored interval. The elapsed times and not-wait times are given in categories of total, mean, and maximum times for each communication line in the Communication-Summary report.

The following example illustrates this report.

You must match which physical devices are using the line to the Stage 1 output from system definition. The line numbers are assigned sequentially, according to their physical occurrence in the Stage 1 input deck.

If your online system specifies the prefetch option for MFS blocks in the control region JCL, the last line of the report contains the statistics for all prefetch events.

|             |                                 |                        |      |         |                      |                     |                     |         |              |
|-------------|---------------------------------|------------------------|------|---------|----------------------|---------------------|---------------------|---------|--------------|
| IMS MONITOR | *****COMMUNICATION SUMMARY***** |                        |      |         | TRACE START 1993 130 | 5:55:15             | TRACE STOP 1993 130 | 5:59:49 | PAGE 0089    |
|             |                                 |                        | (A)  |         |                      | (B)                 |                     |         |              |
| NODE OR     |                                 | .....ELAPSED TIME..... |      |         | NOT IWAIT            | TIME(ELAPSED-IWAIT) |                     |         |              |
| LINE NUMBER | OCCURRENCES                     | TOTAL                  | MEAN | MAXIMUM | TOTAL                | MEAN                | MAXIMUM             |         | DISTRIBUTION |
|             |                                 |                        |      |         |                      |                     |                     |         | NUMBER       |
| PMT01A      | 3                               | 2396                   | 798  | 1547    | 2396                 | 798                 | 1547                |         | 1467A,B      |
| 19          | 182                             | 92155                  | 506  | 1106    | 92155                | 506                 | 1106                |         | 1493A,B      |
| 2           | 59                              | 2280                   | 38   | 41      | 2280                 | 38                  | 41                  |         | 1515A,B      |
| TOTAL       |                                 |                        |      |         |                      |                     |                     |         |              |
|             | 244                             | 96831                  | 396  |         | 96831                | 396                 |                     |         |              |

You can also investigate the amount of data transmitted across nodes with the Line-Functions report. The following example illustrates this report. The report distinguishes between input data and output data. The number of blocks of data and the average and maximum size of the blocks are recorded for data received by IMS and for transmitted data.

This report also includes a measure of how inactive the lines are. An inactive interval is assumed to be the difference between the time that marks the end of the last input block received and the starting time for output transmission. These occurrences of inactivity are termed turnaround intervals, and the report cumulates the number of occurrences as well as the average and maximum times associated with these intervals.

If link traffic for coupled multiple systems is recorded, a set of three reports follows the Line Functions report. These are described in “IMS Monitor reports for MSC” on page 763.

| IMS MONITOR            |                | ****LINE FUNCTIONS*** |                    |                    | TRACE START 1993 130 5:55:15 |                   |                   |                 | TRACE STOP 1993 130 5:59:49 |                  |                  |                | PAGE 0091          |
|------------------------|----------------|-----------------------|--------------------|--------------------|------------------------------|-------------------|-------------------|-----------------|-----------------------------|------------------|------------------|----------------|--------------------|
|                        |                | (A).....              |                    |                    | (B).....                     |                   |                   |                 |                             |                  |                  |                |                    |
| NODE OR<br>LINE NUMBER | DEVICE<br>TYPE | RECEIVE<br>BLOCKS     | RECEIVE<br>BLKSIZE | RECEIVE<br>BLKSIZE | TRANS.<br>BLOCKS             | TRANS.<br>BLKSIZE | TRANS.<br>BLKSIZE | DIST.<br>NUMBER | TURN<br>AROUND<br>INTERVALS | MEAN<br>INTERVAL | MAX.<br>INTERVAL | DIST.<br>NUMB. | PAGING<br>REQUESTS |
| PMT01A                 | 3270V          | 1                     | 29                 | 29                 | 2                            | 170               | 171               | 1468A,B         | 3                           | 798              | 1547             | 1466           | 0                  |
| 19                     | XXXX           |                       |                    |                    |                              |                   |                   |                 | 182                         | 506              | 1106             | 1492           | 0                  |
| 2                      | LOC SYS        |                       |                    |                    |                              |                   |                   |                 | 59                          | 38               | 41               | 1514           | 0                  |
| ALL LINES              |                |                       |                    |                    |                              |                   |                   |                 |                             |                  |                  |                |                    |
|                        |                | 1                     | 29                 |                    | 2                            | 170               |                   |                 | 244                         | 396              |                  |                | 0                  |

The IMS Monitor produces both summary and detailed information on asynchronous processing that takes place in the IMS control region. Asynchronous processing occurs when data transmitted from VTAM arrives. Application program responses also result in asynchronous processing.

This report is similar to the Communication-Summary report because the line number identifies the series of communication processing tasks.

|                              |             |                              |      |                      |          |         |         |            |  |          |     |         |  |           |
|------------------------------|-------------|------------------------------|------|----------------------|----------|---------|---------|------------|--|----------|-----|---------|--|-----------|
| IMS MONITOR                  |             | ****COMMUNICATION IWAIT***** |      | TRACE START 1993 130 |          | 5:55:15 |         | TRACE STOP |  | 1993 130 |     | 5:59:49 |  | PAGE 0090 |
| NODE OR<br>LINE NUMBER       | OCCURRENCES | .....IWAIT TIME.....         |      |                      |          |         |         |            |  | DIST.    |     |         |  |           |
|                              |             | TOTAL                        | MEAN | MAXIMUM              | FUNCTION |         | BLKSIZE | MODULE     |  |          | NO. |         |  |           |
| <hr/>                        |             |                              |      |                      |          |         |         |            |  |          |     |         |  |           |
| ALL LINES...<br>PREFETCH I/O |             |                              |      |                      |          |         |         |            |  |          |     |         |  |           |
| NONE                         |             |                              |      |                      |          |         |         |            |  |          |     |         |  |           |

There are several summary reports that you can use to examine the level of internal contention for resources.

## Pool space failure

The format of the report is shown in the following example.

### POOL SPACE FAILURE SUMMARY

| <i>POOL ID</i> | <i>BYTES REQ.</i> | <i>OCCURRENCES</i> |
|----------------|-------------------|--------------------|
| DLMP           | 8888              | 1                  |
| DLDP           | 7777              | 1                  |
| TOTAL          |                   | 2                  |

### Programs experiencing deadlock

Each time a pair of programs reaches a deadlock over ownership of a segment in a given database data set, the Deadlock-Event-Summary report records the occurrence. Each line in the report shows the two PSBs involved and indicates which is given processing right-of-way (REQ-ING PSB) and which has to reprocess after dynamic backout occurs (LOSING PSB). The report is illustrated in the following example.

#### DEADLOCK EVENT SUMMARY

| REQ-ING PSB | LOSING PSB | DMBNAME  | OCCURRENCES |
|-------------|------------|----------|-------------|
| PSBNAME1    | TPPSBRE3   | DBASEBAL | 1           |
| TOTAL       |            |          | 1           |

### IMS latch conflict

The basic serialization of the task processing in IMS is controlled by ownership of an IMS latch. When different programs are executing, they compete for the ownership. If they wait for the resource, the one possessing the latch has to post the other ITASK waiting for it. Use the Latch Conflict Statistics report to judge the level of contention for a resource.

The latch names, which are abbreviations for the different types of resources being serialized, are as follows:

#### Abbreviation

##### Latch Name

|             |                           |
|-------------|---------------------------|
| <b>DISP</b> | SYS/DISPATCHER            |
| <b>DCSL</b> | DC/CHECKPOINT DC SYSTEM   |
| <b>LUML</b> | DC/LU 6.2 LUM             |
| <b>CONV</b> | DC/CONVERSATION CHECKPT   |
| <b>TERM</b> | DC/TERMINAL               |
| <b>LUBT</b> | DC/LU62 LUB-TIB CHAIN     |
| <b>SCHD</b> | TM/SCHEDULING             |
| <b>TCTB</b> | TM/TCT BLOCK              |
| <b>APSB</b> | TM/ALLOCATE PSB (BLK MVR) |
| <b>PDRB</b> | TM/PDIR BLOCK (BLK MVR)   |
| <b>PSBP</b> | TM/PSB POOL (BLK MVR)     |
| <b>DMBP</b> | TM/DMB POOL (BLK MVR)     |
| <b>PSBB</b> | TM/PSB BLOCK (BLK MVR)    |
| <b>DMBB</b> | TM/DMB BLOCK (BLK MVR)    |
| <b>PDRP</b> | TM/PDIR POOL (BLK MVR)    |
| <b>DBAU</b> | TM/DBRC AUTH (BLK MVR)    |
| <b>DDRB</b> | TM/DDIR BLOCK (BLK MVR)   |
| <b>DDRP</b> | TM/DDIR POOL (BLK MVR)    |

**DBBP** DB/OSAM BUFFER POOL  
**DBLR** DB/DFSDBLR0 MODULE  
**SUBQ** TM/TM SUBQUEUES  
**DBSL** DB/DB CHECKPOINT  
**USER** DC/USER  
**DBLT** RSR SHARING SERIALIZE  
**CCTL** SYS/DBCTL RESOURCE  
**VTCT** SYS/CBTS VTCB POOL  
**VLQB** SYS/CBTS LQB POOL  
**CBTS** SYS/CBTS POOLS (ALL)  
**BLKM**  
     TM/SMB QUEUE HASH TABLE  
**QMGR**  
     SYS/QUEUE MANAGER  
**QBSL** SYS/QUEUE BUFFER  
**SMGT**  
     SYS/STORAGE MANAGEMENT  
**DBLK** SYS/DEPENDENT REGION  
**XCNQ**  
     DB/EXCLUSIVE ENQ/DEQ  
**ACTL** SYS/STATISTICS  
**LOGL** SYS/LOGGER

When a system checkpoint is taken during the time the monitor is active, latch conflict statistics are reset to zero, thus corrupting the values presented in this report. If this situation exists, the following message will be inserted at the top of the report:

```

**** A CHECKPOINT OCCURRED DURING MONITOR RUN ****
**** LATCH CONFLICT STATISTICS ARE INVALID ****
**** SEE UTILITIES REFERENCE MANUAL ****
    
```

However, if the master terminal operator issues the /CHECKPOINT command with the STATISTICS keyword parameter, latch conflict statistics are reset to zero, but the IMS monitor is not notified. Therefore, DFSUTR20 cannot detect that the statistics have been corrupted and does not issue this message.

**Recommendation:** Do not issue statistics checkpoints while the monitor is running.

The different types of latches and the counters that exhibit the level of contention are given in the Latch Conflict Statistics report. The following example shows a sample of the latch-conflict-statistics report. The entries are organized according to the latch names.

```

            IMS MONITOR   ** GENERAL REPORTS **           TRACE START 1993 209...

                        LATCH CONFLICT STATISTICS

LATCH      COUNT              AT              AT
    
```

| <i>NAMES</i> | <i>FIELD</i> | <i>START</i> | <i>END</i> | <i>DIFF.</i> |
|--------------|--------------|--------------|------------|--------------|
| LOGL         | CONTENTIONS  | 0            | 0          | 0            |
| SMGT         | CONTENTIONS  | 0            | 0          | 0            |
| XCNQ         | CONTENTIONS  | 0            | 0          | 0            |
| ACTL         | CONTENTIONS  | 0            | 0          | 0            |
| CBTS         | CONTENTIONS  | 0            | 0          | 0            |
| DBLK         | CONTENTIONS  | 0            | 0          | 0            |

## Using frequency distributions from IMS Monitor output

The reports derived from the IMS Monitor data records contain many summary lines where the mean time is given. If you are interested in the distribution of those timed events, rather than just average and maximum times, you can request the Report Print utility to individually record the events in a frequency distribution across a range of intervals.

Some distributions are not time dependent, such as those for transaction queue loads or transmitted block sizes.

The following tables show the major IMS Monitor reports and the type of frequency distributions generated for each report. Each type results in several distributions, depending on how many entries are in each section of the report. For each type of frequency distribution, the data is cumulated in suitable intervals or ranges. The set of ranges used for each type is given an identifier, shown in the ID column.

- The following table shows the report distributions sorted by region summary.

*Table 60. Distribution reports by region summary*

| <b>Report name</b>            | <b>ID</b> | <b>Description</b> |
|-------------------------------|-----------|--------------------|
| Schedule end to 1st DL/I call | D1        | Elapsed time       |
| Elapsed execution time        | D2        | Not wait time      |
| DL/I calls                    | D3        | N/A                |
|                               | D4        | N/A                |
| External Subsystem calls      | D5        | Elapsed time       |
| Waits per DL/I call           | D6        | Not wait time      |
| Idle for intent               | D43       | Elapsed time       |
| Checkpoint                    | D7        | N/A                |
|                               | D8        | N/A                |
|                               | D20       | Elapsed time       |
|                               | D21       | Not wait time      |

- The following table shows the report distributions sorted by program region.

*Table 61. Report distributions by program region*

| <b>Report name</b>            | <b>ID</b> | <b>Description</b> |
|-------------------------------|-----------|--------------------|
| Elapsed execution time        | D30       | N/A                |
| Schedule end to 1st DL/I call | D31       | N/A                |

- The following table shows the report distributions sorted by program summary.

*Table 62. Report distributions by program summary*

| Report name                        | ID  | Description |
|------------------------------------|-----|-------------|
| Processor time per schedule        | D15 | N/A         |
| Transactions dequeued per schedule | D14 | N/A         |
| Elapsed time per schedule          | D9  | N/A         |
| Schedule end to 1st DL/I call      | D10 | N/A         |

- The following table shows the report distributions sorted by communication summary.

*Table 63. Report distributions by communication summary*

| Report name        | ID  | Description |
|--------------------|-----|-------------|
| Line elapsed time  | D18 | N/A         |
| Line not wait time | D19 | N/A         |

- The following table shows the report distributions sorted by line functions.

*Table 64. Report distributions by line functions*

| Report name              | ID  | Description |
|--------------------------|-----|-------------|
| Received block length    | D36 | N/A         |
| Transmitted block length | D37 | N/A         |
| Inactive intervals       | D38 | N/A         |

- The following table shows the report distributions sorted by MSC queuing summary.

*Table 65. Report distributions by MSC queuing summary*

| Report name   | ID  | Description |
|---------------|-----|-------------|
| Time in queue | D39 | N/A         |

- The following table shows the report distributions sorted by transaction queuing.

*Table 66. Report distributions by transaction queuing*

| Report name                        | ID  | Description   |
|------------------------------------|-----|---------------|
| Transactions on queue at schedule  | D17 | N/A           |
| Transactions dequeued per schedule | D16 | N/A           |
| Prefetch format blocks             | D28 | Elapsed time  |
|                                    | D29 | Not wait time |

- The following table shows the report distributions sorted by call summary.

*Table 67. Report distributions by call summary*

| Report name                           | ID  | Description |
|---------------------------------------|-----|-------------|
| PSB waits per DL/I call               | D13 | N/A         |
| PSB waits per external subsystem call | D44 | N/A         |
| PSB elapsed time per call             | D11 | N/A         |

Table 67. Report distributions by call summary (continued)

| Report name                  | ID  | Description  |
|------------------------------|-----|--------------|
| PSB not wait time per call   | D12 | N/A          |
| PSB external subsystem calls | D45 | Elapsed time |

- The following table shows some distributions derived from buffer pool statistics for wait times.

Table 68. Wait time distributions

| Function                 | ID  | Module key |
|--------------------------|-----|------------|
| Storage                  | D22 | SMN        |
| OSAM I/O                 | D23 | DBH        |
| VSAM I/O                 | D24 | VBH        |
| Scheduler internal       | D25 | MSC        |
| Queue manager I/O        | D26 | QMG        |
| Block loader I/O         | D27 | BLR        |
| MFS block I/O            | D32 | MFS        |
| MFS directory I/O        | D33 | MFS        |
| HSAM I/O                 | D34 | DIE        |
| Format Buffer Pool Space | D35 | PMM        |
| PI enqueue               | D40 | None       |
| QMGR SNAPQ Check         | D42 | None       |

**Related reference:**

“Default values of distribution definitions” on page 761

## How to get a frequency distribution output

To request the IMS Monitor Report Print utility to gather distribution data, you must include a DIS input control statement. This causes all report items with an entry under a column headed MEAN to have a corresponding frequency distribution as part of the Distribution Appendix.

Each report line includes an identifying reference number under the column headed Distribution Number so that you can locate the distribution data in the appendix, flagged by that same number.

## How frequency distribution ranges are defined

A set of ten intervals is defined for each summary line and the occurrences falling in each interval are cumulated. The interval ranges are preset with default end points. The default end points are chosen so that they are suitable to the event.

For example, the end points, for DL/I call elapsed time are: 0, 1000, 2000, 4000, 8000, 16000, 32000, 64000, 128000, 256000, INF (all times are in milliseconds). The lower limit of the first interval always defaults to zero, and the upper limit of the tenth interval is infinity (INF).

Although several types of distribution can use the same set of end points, each type is assigned a distribution identifier. You can use this to redefine the end points. To override the default end points, include an input control statement to

the Report Print utility. The statement specifies the type of distribution identifier and gives the desired end point values. For example, the DL/I call elapsed time end points could be respecified by:

D5 0,500,1000,1500,2000,4000,,,100000,500000

The values of the unspecified end points remain at their default values of 32000 and 64000 as does the last (INF).

The following example, which is a sample page taken from a Distribution Appendix, shows how individual distributions are numbered, and how ranges vary with the type of distribution. The lines are arranged in pairs, with the second one recording the cumulated counts.

```

IMS MONITOR ****DISTRIBUTION APPENDIX**** TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0200
# 1.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 2.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 3.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 4.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 5.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 6.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 7.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 8.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 9.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 10.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 11.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 12.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 13.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 14.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 15.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 16.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF

```

## Default values of distribution definitions

Using an identifier provided in the frequency distribution tables and the Wait Time Distributions table, you can determine the default end points for the distribution by locating it in the following list.

**D1, D2, D5, D6, D9, D10, D11, D12, D15 D18, D19, D20, D21, D22, D25, D27 D28, D29, D30, D31, D43, and D45**

0, 1000, 2000, 4000, 8000, 16000, 32000, 64000, 128000, 256000, INF

**D3** 0, 50000, 100000, 150000, 200000, 250000, 300000, 350000, 400000, 450000, INF

**D4** 0, 200000, 400000, 600000, 800000, 1000000, 1200000, 1400000, 1600000, 1800000, INF

**D7, D13, and D44**

0, 0, 1, 2, 3, 4, 5, 6, 7, 8, INF

**D8** 0, 100000, 200000, 300000, 400000, 500000, 600000, 700000, 800000, 900000, INF



**D14, D16, D17**

0, 1, 2, 3, 4, 5, 10, 15, 30, 90, INF

**D23, D24, D26, D32, D40, D42**

0, 2000, 8000, 24000, 50000, 100000, 150000, 200000, 250000, 300000, INF

**D33, D34, D35**

0, 2000, 4000, 8000, 16000, 32000, 64000, 96000, 128000, 160000, INF

**D36, D37**

0, 10, 20, 40, 80, 100, 200, 400, 800, 1000, INF

**D38** 0, 1000, 10000, 100000, 200000, 500000, 800000, 1000000, 1500000, 2000000, INF

**D39** 0, 1000, 5000, 10000, 50000, 100000, 500000, 1000000, 5000000, 10000000, INF

**Related concepts:**

“Using frequency distributions from IMS Monitor output” on page 758

## Interpreting the Distribution Appendix

You can use the detailed output in the Distribution Appendix when you suspect an unusual combination of events was reported in a report summary line. Usually, the average and maximum times or counts are sufficient to highlight a resource usage problem.

However, if you suspect the mean value to be masking an unusual distribution you can draw on the detail contained in the IMS Monitor output records.

For example, suppose you are investigating a change in the scheduling algorithm for a particular transaction and need to know how many transactions were able to be processed for each scheduling of an application program. The following figure shows a possible histogram for the processed transactions:




Figure 89. Number of transactions processed for each scheduling of an application program

The Average=2.5 transactions pre schedule. The distribution in the figure suggests that many schedules were able to process only one or two transactions, and few schedules significantly exhausted the queue. The distribution data for the histogram is as follows:

|                       |   |   |   |   |   |      |     |
|-----------------------|---|---|---|---|---|------|-----|
| Number of schedules   | 1 | 2 | 3 | 4 | 5 | 6-10 | >10 |
| Transactions dequeued | 8 | 7 | 0 | 0 | 2 | 1    | 0   |

The Distribution Appendix presents the histogram data in the form of two lines:

- The first line shows the intervals, prefixed by a cross reference to an individual line on the earlier output.
- The second line gives the number of events occurring in those intervals.

This data appears as follows:

```
# 950B...0...1...2...3...4...5...10...15...30...90...INF
      8   7   0   0   2   1   0   0   0   0
```

The cross reference 950B points to a unique report line. For example, the Transaction-Queuing report on the appropriate line for the transaction of interest show, 950A,B under the column headed DISTRIBUTION NUMBER. Use the reference number 950B to locate the data in the Distribution Appendix. The 950A reference points to the data for the number of transactions in the queue at schedule time.

---

## IMS Monitor reports for MSC

The IMS Monitor Report Print program includes three reports that highlight message events caused by Multiple Systems Coupling.

- MSC Traffic report

This report shows the enqueue and dequeue counts of messages that use the various link paths for the monitor interval.

- MSC Summary report

This report shows summaries of:

- The traffic queues for each input transaction name
- The traffic queues for each destination name
- The traffic queues for each link number
- The traffic queues for each destination system

- MSC Queuing Summary report

This report is generated only when intersystem messages are queued on the local system before being sent to the destination system. The local system must be an intermediate system. This report shows:

- Maximum time messages spend in queues
- Average time messages spend in queues
- Maximum queue lengths
- Maximum queue counts
- Total number of messages queued for all links the local system participates in

All three of the reports can have entries in the Distribution Appendix. You can examine the frequency distributions of the traffic if you suspect unusual transmission patterns.

**Related concepts:**

“Overview of IMS Monitor reports” on page 788

“Monitoring line activity” on page 803

#### Related tasks:

 Monitoring and tuning multiple systems (Communications and Connections)

## MSC Traffic report

The MSC Traffic report reveals the individual queue loads for all traffic between partner systems of which the monitored system is the local system. The report lists all the unique system identification numbers (SIDs) that are defined for communications for that local system.

It then summarizes the total messages queued and dequeued for each combination of the following variables:

- Input name (terminal or program that was a message source)
- Destination name (terminal or program)
- Input system (SID)
- Destination system (SID)
- Link number
- Link type (CTC, MTM, or VTAM)

The following example illustrates the MSC-traffic report. If a message originates in the local system, its presence is accounted for in the dequeue counts only.

Messages with local destinations appear only in the enqueue count.

```
IMS MONITOR *****MSC TRAFFIC REPORT***** TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0151
LOCAL SID VALUES = 101, 102, 103, 104, 105, 106, 107, 108, 109, 110, 111, 112, 113, 114, 115
```

| INPUT<br>NAME | DESTIN.<br>NAME | INPUT<br>SID | DEST.<br>SID | LINK<br>NO. | LINK<br>TYPE | ENQUEUE<br>COUNT | DEQUEUE<br>COUNT |
|---------------|-----------------|--------------|--------------|-------------|--------------|------------------|------------------|
| DSWT3685      | DSWT3685        | 2            | 2            | 2           | C-C          | 0                | 1                |
|               | SC6Y            | 2            | 102          | 2           | C-C          | 1                | 0                |
| DSWT6161      | DSWT6161        | 3            | 3            | 3           | C-C          | 0                | 1                |
|               | SC6Z            | 3            | 103          | 3           | C-C          | 1                | 0                |
| DSWT3838      | DSWT3838        | 2            | 2            | 2           | C-C          | 0                | 1                |
|               | SC6Y            | 2            | 102          | 2           | C-C          | 1                | 0                |
| DSWT4618      | DSWT4618        | 3            | 3            | 3           | C-C          | 0                | 1                |
|               | SC6Z            | 3            | 103          | 3           | C-C          | 1                | 0                |
| DSWT3903      | DSWT3903        | 3            | 3            | 3           | C-C          | 0                | 1                |
|               | SC2Z            | 3            | 103          | 3           | C-C          | 1                | 0                |
| DSWT5418      | DSWT5418        | 3            | 3            | 3           | C-C          | 0                | 1                |
|               | SC4Z            | 3            | 103          | 3           | C-C          | 1                | 0                |
| DSWT4673      | DSWT4673        | 3            | 3            | 3           | C-C          | 0                | 1                |
|               | SC2Z            | 3            | 103          | 3           | C-C          | 1                | 0                |
| DSWT5141      | DSWT5141        | 45           | 45           | 45          | VTAM         | 0                | 1                |
|               | PS3X            | 45           | 145          | 45          | VTAM         | 1                | 0                |
| DSWT4391      | DSWT4391        | 2            | 2            | 2           | C-C          | 0                | 1                |
|               | SC4Y            | 2            | 102          | 2           | C-C          | 1                | 0                |
| DSWT3324      | DSWT3324        | 17           | 17           | 17          | VTAM         | 0                | 1                |
|               | IT1Y            | 17           | 117          | 17          | VTAM         | 1                | 0                |
| DSWT4781      | DSWT4781        | 3            | 3            | 3           | C-C          | 0                | 1                |
|               | SC4Z            | 3            | 103          | 3           | C-C          | 1                | 0                |
| DSWT3525      | DSWT3525        | 3            | 3            | 3           | C-C          | 0                | 1                |
|               | SC6Z            | 3            | 103          | 3           | C-C          | 1                | 0                |
| DSWT4542      | DSWT4542        | 2            | 2            | 2           | C-C          | 0                | 1                |
|               | SC6Y            | 2            | 102          | 2           | C-C          | 1                | 0                |
| DSWT5796      | DSWT5796        | 3            | 3            | 3           | C-C          | 0                | 1                |
|               | SC2Z            | 3            | 103          | 3           | C-C          | 1                | 0                |
| DSWT4782      | DSWT4782        | 3            | 3            | 3           | C-C          | 0                | 1                |
|               | SC6Z            | 3            | 103          | 3           | C-C          | 1                | 0                |
| DSWT4633      | DSWT4633        | 2            | 2            | 2           | C-C          | 0                | 1                |
|               | SC6Y            | 2            | 102          | 2           | C-C          | 1                | 0                |
| DSWT3655      | DSWT3655        | 12           | 12           | 12          | VTAM         | 0                | 1                |
|               | SC6U            | 12           | 112          | 12          | VTAM         | 1                | 0                |
| DSWT3892      | DSWT3892        | 2            | 2            | 2           | C-C          | 0                | 1                |
|               | SC6Y            | 2            | 102          | 2           | C-C          | 1                | 0                |
| DSWT3338      | DSWT3338        | 4            | 4            | 4           | VTAM         | 0                | 1                |

|               |          |   |     |   |      |      |      |
|---------------|----------|---|-----|---|------|------|------|
|               | SC2U     | 4 | 104 | 4 | VTAM | 1    | 0    |
| DSWT4681      | DSWT4681 | 3 | 3   | 3 | C-C  | 0    | 1    |
|               | SC4Z     | 3 | 103 | 3 | C-C  | 1    | 0    |
| DSWT4482      | DSWT4482 | 2 | 2   | 2 | C-C  | 0    | 1    |
|               | SC6Y     | 2 | 102 | 2 | C-C  | 1    | 0    |
| DSWT4902      | DSWT4902 | 3 | 3   | 3 | C-C  | 0    | 1    |
|               | SC2Z     | 3 | 103 | 3 | C-C  | 1    | 0    |
| DSWT4558      | DSWT4558 | 3 | 3   | 3 | C-C  | 0    | 1    |
| DSWT4558      | SC6Z     | 3 | 103 | 3 | C-C  | 1    | 0    |
| DSWT3925      | DSWT3925 | 2 | 2   | 2 | C-C  | 0    | 1    |
| TOTAL TRAFFIC |          |   |     |   |      | 1353 | 1359 |

#### Related tasks:

 Monitoring and tuning multiple systems (Communications and Connections)

## MSC Summaries report

The MSC Summaries report shows you the enqueue and dequeue activity for messages that are handled by the local system but are part of the Multiple System Coupling (MSC) traffic.

The report format is illustrated in the following example.

The first set of queuing counts shows how many transactions of each type were entered in the monitor interval, and how many were subsequently dequeued.

The second set of counts summarizes the total traffic for each destination name. You can distinguish the primary transactions and responses by the resource names and examine the relative servicing of the link transmissions using the difference between the enqueue and dequeue counts.

The third set of counts lists the active links by link number, and you can determine if there is buildup on the link by the difference in the enqueue and dequeue counts.

The fourth set of counts records the traffic that is going to other systems by all link paths. You can judge by the difference in enqueue and dequeue counts whether the overall pattern of link priorities to one or more systems is causing buildup of cross-system traffic.

```

IMS MONITOR *****MSC SUMMARIES***** TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0178
<--SUMMARY BY INPUT NAME--> <--SUMMARY BY DESTINATION NAME--> <--SUMMARY BY LOGICAL LINK--> <--SUMMARY BY DEST. SYS. ID-->

```

| INPUT<br>NAME | ENQUEUE<br>COUNT | DEQUEUE<br>COUNT | DESTIN.<br>NAME | ENQUEUE<br>COUNT | DEQUEUE<br>COUNT | LINK<br>NO. | ENQUEUE<br>COUNT | DEQUEUE<br>COUNT | DEST.<br>SID | ENQUEUE<br>COUNT | DEQUEUE<br>COUNT |
|---------------|------------------|------------------|-----------------|------------------|------------------|-------------|------------------|------------------|--------------|------------------|------------------|
| DSWT3577      | 1                | 1                | DSWT4358        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT4048      | 1                | 1                | DSWT5988        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT5216      | 1                | 1                | DSWT5457        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT4776      | 1                | 1                | DSWT5187        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT5496      | 1                | 1                | DSWT3312        | 0                | 2                |             |                  |                  |              |                  |                  |
| DSWT5277      | 1                | 1                | DSWT5604        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT5711      | 1                | 1                | DSWT3347        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT5274      | 1                | 1                | DSWT5338        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT5807      | 1                | 1                | DSWT3268        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT3685      | 1                | 1                | DSWT3676        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT6161      | 1                | 1                | DSWT5428        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT3838      | 1                | 1                | DSWT5395        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT4618      | 1                | 1                | DSWT4168        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT3903      | 1                | 1                | DSWT5061        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT5418      | 1                | 1                | DSWT3511        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT4673      | 1                | 1                | DSWT3363        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT5141      | 1                | 1                | DSWT3674        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT4391      | 1                | 1                | DSWT4467        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT3324      | 1                | 1                | DSWT4501        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT4781      | 1                | 1                | DSWT5037        | 0                | 1                |             |                  |                  |              |                  |                  |
| DSWT3525      | 1                | 1                | DSWT4298        | 0                | 1                |             |                  |                  |              |                  |                  |

|          |   |   |          |   |   |
|----------|---|---|----------|---|---|
| DSWT4542 | 1 | 1 | DSWT5778 | 0 | 1 |
| DSWT5796 | 1 | 1 | DSWT4003 | 0 | 1 |
| DSWT4782 | 1 | 1 | DSWT3988 | 0 | 1 |
| DSWT4633 | 1 | 1 | DSWT4217 | 0 | 1 |
| DSWT3655 | 1 | 1 | DSWT6135 | 0 | 1 |
| DSWT3892 | 1 | 1 | DSWT5147 | 0 | 1 |
| DSWT3338 | 1 | 1 | DSWT5381 | 0 | 1 |
| DSWT4681 | 1 | 1 | DSWT5593 | 0 | 1 |
| DSWT4482 | 1 | 1 | DSWT3304 | 0 | 1 |
| DSWT4902 | 1 | 1 | DSWT5081 | 0 | 1 |
| DSWT4558 | 1 | 1 | DSWT4671 | 0 | 1 |
|          |   |   | DSWT3655 | 0 | 1 |
|          |   |   | DSWT3892 | 0 | 1 |
|          |   |   | DSWT3338 | 0 | 1 |
|          |   |   | DSWT4681 | 0 | 1 |
|          |   |   | DSWT4482 | 0 | 1 |
|          |   |   | DSWT4902 | 0 | 1 |
|          |   |   | DSWT4558 | 0 | 1 |
|          |   |   | DSWT3925 | 0 | 1 |

**Related tasks:**

 Monitoring and tuning multiple systems (Communications and Connections)

**MSC Queuing Summary report**

The MSC Queuing Summary report provides information about intersystem message traffic only. You can use the sample of traffic recorded in the IMS Monitor interval to examine the maximum and average time messages spend in queues waiting to be sent on active links.

You can detect whether the link priorities are causing undue delay of primary messages through the intermediate system, or whether there is a build up of responses. The report shows the logical link paths for this system which is an intermediate system. Each incoming link number shows the number of messages that are queued before transmission on their specified outward bound link number. The maximum queue count is given as well as the maximum and average time on the intermediate system queues.

The following example shows the MCS queuing summary report.

|                                                              |      |            |      |          |        |        |        |  |
|--------------------------------------------------------------|------|------------|------|----------|--------|--------|--------|--|
| IMS MONITOR *****MSC QUEUING SUMMARY*****                    |      |            |      |          |        |        |        |  |
| PAGE 0022                                                    |      |            |      |          |        |        |        |  |
| TRACE START 2011 063, 10:33:13 TRACE STOP 2011 063, 10:34:54 |      |            |      |          |        |        |        |  |
| ENQUE.....                                                   |      | DEQUE..... |      |          | MAX.Q  | MAX.   | MEAN   |  |
| LINK NO.                                                     | TYPE | LINK NO.   | TYPE | MESSAGES | LENGTH | Q TIME | Q TIME |  |
| <hr/>                                                        |      |            |      |          |        |        |        |  |
| 13                                                           | VTAM | 6          | M-M  | 2        | 1      | 2342   | 1234   |  |
| TOTALS...                                                    |      |            |      | 2        |        |        | 1234   |  |

**Related tasks:**

 Monitoring and tuning multiple systems (Communications and Connections)

---

## Chapter 49. IMS Monitor reports for DBCTL

DBCTL monitoring provides data about the processing that occurs when a CCTL transaction accesses DBCTL databases. The CCTL gains this access using database resource adapter (DRA) requests.

This topic describes:

- The events that the IMS Monitor collects
- The content of the reports produced by the IMS Monitor Report Print Program

Monitoring has different meanings for DBCTL and DB/DC. For DB/DC, the end user enters the transaction on a terminal. The transaction is processed by IMS and then returns a result to the user. Transaction characteristics that are monitored include total response time and the occurrences of resource contentions (for example, PSB schedule wait time, and database I/Os).

DBCTL has neither transactions nor terminal end users. It does, however, work on behalf of transactions entered by CCTL terminal users.

A typical sequence of these DRA requests would be:

1. A SCHED request to get a PSB scheduled in DBCTL
2. A DL/I request to make database calls
3. A sync-point request, COMMTERM, to commit the updates and release the PSB

The DBCTL process that encompasses these requests is called a unit of recovery (UOR).

DBCTL provides monitoring data about UORs, such as: total time UOR existed, wait time for PSB schedule, and I/Os during database calls. This information is very similar to IMS transaction monitor data. In a DBCTL-CCTL system, however, the UOR data represents only part of the total processing of a CCTL transaction. Therefore, CCTL monitor data is necessary to get a total view of CCTL transaction performance.

DBCTL does not change the format or usage of the IMS monitor reports. There are reports and fields within reports that are not applicable to DBCTL. Generally, these are in the transaction manager and communication areas. There are some fields that are interpreted differently in a DBCTL environment.

For reports that do not apply to DBCTL, either a heading without data is shown or no report is generated. These reports are:

- Message Queue Pool report
- Message Format Buffer Pool report
- Communication Summary report
- Communication IWAIT report
- Line Functions report
- MSC Traffic report
- MSC Summaries report
- MSC Queuing Summary report

The term *region* in IMS Monitor reports refers to a PST assigned to a specific dependent region that processes specific IMS transactions. In DBCTL monitor reports the term *region* still applies to a PST. A PST can service one CCTL thread (transaction) at a time. However, CCTL threads change, resulting in one PST servicing many different CCTL transactions. Since multiple CCTLs can connect to DBCTL, the PST can actually service transactions from different CCTLs.

All of the threads built for a CCTL carry the job name of the CCTL. This appears as the same job name for many regions in the General Reports.

Within a trace interval, a thread can be assigned to multiple CCTLs, but it can only be assigned to one CCTL at any instant of time. So, depending on the number of CCTLs attached to DBCTL, the Region Summary reports can show:

- One region with only one job name.
- One region with different job names.
- Multiple regions with different job names. Some regions can have the same job name and some can have different job names.
- Multiple regions with only one job name.

Any monitor report for a region is a summary of all the CCTLs a thread served during the trace interval (for example, the elapsed time for all CCTLs that a thread has been assigned to during the trace interval).

The reports generated by the IMS Monitor are the same for BMPs and non-message BMPs.

UOR elapsed times are spent in DBCTL, not in the DRA. The time spent in the DRA is considered part of the CCTL, therefore the DRA time is not reported by any DBCTL statistics.

---

## IMS Monitor trace event intervals

The IMS Monitor trace interval is defined by the master terminal operator's use of the /TRACE command between the start and stop command entries. The online IMS events are recorded in IMS Monitor records placed in the IMSMON data set.

The event timings are related to dependent region activity. The following figure shows the boundaries of the timed event intervals.

The Monitor trace interval includes the following intervals:

- Scheduling and Termination
  - Block loader busy
  - Intent failures (exclusive intent and data sharing) and Schedule failures (PSB busy and space failure)
  - Sched/Term elapsed
    - NOT-WAIT
    - ACBLIB waits
    - DB Flush waits
    - DB CLOSE waits
- Region occupancy (which overlaps with all of Sched/Term elapsed)
  - Schedule to first call
  - Elapsed execution

The NOT-WAIT time for a region is the elapsed time not accounted for by wait time. Any delay coming from either paging or the processor being dispatched for a higher priority task results in an increase in the NOT-WAIT times.




Figure 90. IMS Monitor trace event intervals

## Overview of IMS Monitor reports

A list of reports available from data collected by the IMS Monitor, together with the principal performance data they contain, is shown in the following table.

**Related reading:** For a description of those reports marked “DB”, see Chapter 47, “DB Monitor reports,” on page 709.

The order of the reports listed in the table matches the sequence of the output from the IMS Monitor Report Print program. The duration of a monitoring snapshot might not include certain events necessary for an individual report, in which case only report headings or partial data are produced.

Table 69. IMS Monitor reports output sequence and information

| Report name               | Principal information                     |
|---------------------------|-------------------------------------------|
| System Configuration      | Monitor run documentation                 |
| Database Buffer Pool (DB) | Count of DB calls and I/O per transaction |
| VSAM Buffer Pool (DB)     | Count of inserts and I/Os                 |
| Latch Conflict Statistics | IMS internal processing                   |
| Region and Jobname        | Monitor run documentation                 |
| Region Summary            | Elapsed times and count of DL/I calls     |
| Region Wait               | Wait times                                |
| Programs by Region        | Elapsed times for region usage            |
| Program Summary           | Overall program statistics                |
| Program I/O (DB)          | Wait times/PCB                            |
| Reports                   | Count of space failures and deadlocks     |
| Run Profile               | Monitor run documentation                 |



Table 69. IMS Monitor reports output sequence and information (continued)

| Report name           | Principal information                |
|-----------------------|--------------------------------------|
| Call Summary (DB)     | Call counts and timings/segment type |
| Distribution Appendix | Event frequency distributions        |

The majority of the data items in IMS Monitor reports are elapsed times. These are normally expressed in microseconds. An entry of 1876534 is 1.876534 seconds or 1876 milliseconds. Any times that do not follow this convention show the unit of measure on the report.

You can also find counts of events under the heading OCCURRENCES, and some figures that represent the number of bytes.

## Documenting the monitoring run

For each trace interval, several general reports or overall summaries are generated for the processing that took place. You can use these reports as part of your IMS Monitor run documentation.

It is important to record, as accurately as possible, the conditions under which the trace was taken. Your documentation can include system status information (obtained by the /DISPLAY command) several times before and after the trace, an expected profile of the CCTL transaction activity, and any desired processing events. The trace interval should represent typical processing loads and not be a biased or inadequate historical record.

### Adding to the System Configuration report data

The first general report (titled SYSTEM CONFIGURATION) is found under the page heading BUFFER POOL STATISTICS. It shows the modification level of the IMS and z/OS systems. You can choose to add a list of IMS APARs applied and include the service levels of the application programs, especially if the latter are not permanent programs or are part of a staged implementation. The system configuration output is illustrated in the following section.

### Recording the Monitor trace interval

The heading of most IMS Monitor reports carries the trace start and stop times. It is shown in the format YEAR DAY (Julian) HH:MM:SS. The overall length of the trace interval is given in milliseconds under the title MONITOR OVERHEAD DATA. The following example shows how many trace records were placed on the IMSMON data set. An example of the monitor trace interval recording is shown in the following figure.

```
***IMS  MONITOR***  BUFFER POOL STATISTICS  TRACE START 1993 130  5:55:15  TRACE STOP 1993 130  5:59:49  PAGE 0001
S Y S T E M   C O N F I G U R A T I O N
```

```
SYSTEM CONFIGURATION      :
  IMS VERSION              :    4
  RELEASE LEVEL            :
  MODIFICATION NUMBER      :
```

## Completing the Monitor run profile

A compact set of processing ratios will be found at the end of the Run Profile report. The statistics summarize, for the monitor interval, the UOR throughput and the degree of DL/I and I/O activity. An example of the report is shown in the following example.

```
IMS MONITOR  **RUN PROFILE**          TRACE START 1993 130   5:55:15   TRACE STOP 1993 130   5:59:49   PAGE 0184
TRACE ELAPSED TIME IN SECONDS.....274.6
TOTAL NUMBER OF MESSAGES DEQUEUED.....1403
TOTAL NUMBER OF SCHEDULES.....173
NUMBER OF TRANSACTIONS PER SECOND.....5.1
TOTAL NUMBER OF DL/I CALLS ISSUED.....18632
NUMBER OF DL/I CALLS PER TRANSACTION.....13.2
NUMBER OF OSAM BUFFER POOL I/O'S.....11236,      8.0 PER TRANSACTION
NUMBER OF MESSAGE QUEUE POOL I/O'S.....0,         0.0 PER TRANSACTION
NUMBER OF FORMAT BUFFER POOL I/O'S.....0,         0.0 PER TRANSACTION
RATIO OF PROGRAM ELAPSED TO DL/I ELAPSED:
    REGION 1: 1.09
    REGION 2: 1.09
    REGION 3: 1.00
    REGION 4: 1.02
    REGION 5: 1.01
    REGION 6: 1.00
    REGION 7: 1.00
    REGION 8: 1.00
    REGION 9: 1.17
    REGION 10: 1.00
    REGION 11: 1.00
    REGION 49: 1.03
    REGION 50: 1.19
RATIO OF DL/I ELAPSED TO DL/I IWAIT:
    REGION 1: 325.65
    REGION 2: 73.49
    REGION 4: 100.35
    REGION 5: 85.76
    REGION 6: 82.99
    REGION 47: 95.64
    REGION 48: 45.93
    REGION 49: 9.22
```

The lower half of the Run Profile report shows several ratios:

- Program elapsed time to DL/I elapsed time for each region
- DL/I elapsed time to wait time during DL/I processing
- Program elapsed time to other subsystem call elapsed time
- DL/I elapsed time to other subsystem call elapsed time

Each region is identified by a sequence number, starting at region 1.

There are some generalized processing ratios that are given at the end of several buffer pool statistics reports. You can include them in the documented profile of the trace interval. These are not specific to one UOR or system resource but can be used as indicators of variation across a series of monitor runs. The ratios are:

- The total number of OSAM reads + OSAM writes + all waits divided by the total number of transactions.

From the Message Queue Pool report, this ratio indicates on a per transaction basis the physical I/O activity required to handle the scheduling function.

- The total number of OSAM reads + OSAM writes + BISAM reads divided by the total number of transactions.

From the Database Buffer Pool report, this ratio indicates on a per transaction basis the physical I/O activity required to handle the database buffering function.

## Verifying IMS Monitor report occurrences

When you examine the output from the IMS Monitor Report Print Program, the presence of a report heading does not necessarily mean that appropriate data will be listed. System definition options and utility control statements also affect the content of the output as follows:

- The output does not include a Call Summary report unless a control statement specifies DLI.
- The output does not include a set of Distribution reports unless a control statement specifies DIS or DISTRIBUTION. The column headed DISTRIBUTION NUMBER that occurs on many of the reports contains cross-references to items included in the Distribution reports.
- The output consists of just a Call Summary report if a control statement specified ONLY DLI.

Because many of the summary reports require system status to calculate the difference between start and end values, and this status is obtained during the /TRACE SET OFF processing, the IMS Monitor execution must end before termination of the IMS control region. If the trace was not stopped properly, the following message is issued:

```
NO DATABASE BUFFER POOL TRACES AT END TIME ON MONITOR LOG TAPE
****DATABASE BUFFER POOL REPORT CANCELLED****
```

Similarly, other summary reports are not produced.

The series of reports titled Buffer Pool Statistics do not include a VSAM Buffer Pool section unless the database in IMS.ACBLIB uses the VSAM access method. If VSAM is not used, the following message is issued:

```
NO VSAM BUFFER POOL TRACES ON MONITOR LOG TAPE
****VSAM BUFFER POOL REPORT CANCELLED****
```

If the source data used to formulate a particular IMS Monitor report, or a section of that report, has not been recorded by the IMS Monitor during the trace interval, the report contains only the headings.

---

## Monitoring activity in dependent regions

The IMS Monitor gathers timing information for every dependent region identified in the /TRACE command (a CCTL thread) active during the trace interval. It records the total of the elapsed times for each event, the maximum individual time encountered, and the average time.

There are three major reports that display timings. The reports and a list of their content are:

- **Region Summary Report**
  - Scheduling and termination
  - Schedule end to first call
  - Elapsed execution with separate summaries shown for:
    - DL/I calls
    - External subsystem service and command calls
    - External subsystem database access calls
    - Checkpoint processing
    - Region occupancy

- **Region Wait**
  - Waits during scheduling and termination
  - Waits during DL/I calls
  - Waits during external subsystem calls
  - Waits during checkpoint
- **Programs by Region**
  - Elapsed execution
  - Schedule end to first call

In this report, “program name” is the PSB name for the UOR.

These three reports are illustrated in the following examples.

Activities in dependent regions are placed in five timing categories:

- Elapsed time for scheduling and termination
 

The scheduling process includes many preparatory events such as block loading from an active IMS.ACBLIBA/B data set, and obtaining ownership of the PSB. The time required to terminate is the time it takes DBCTL to complete this process after receiving a request to terminate the UOR.
- Elapsed time from end of schedule to first call
 

This is the time from when DBCTL completes scheduling until the time DBCTL reviews the first DL/I call. Events that occur during this time are all outside of DBCTL, either in the database resource adapter (DRA) or the CCTL.
- Program elapsed time, including all calls
 

This time encompasses the major UOR processing, measured from the first DL/I call to the call that terminates a UOR.
- Elapsed time performing DL/I calls
 

This time includes all DL/I calls. The time in DBCTL is recorded and summed.

The following example shows a region summary report.

| IMS MONITOR ****REGION SUMMARY**** |    |     | TRACE START 1993 130 5:55:15 |           |           | TRACE STOP 1993 130 5:59:49   |      |             | PAGE 0011    |
|------------------------------------|----|-----|------------------------------|-----------|-----------|-------------------------------|------|-------------|--------------|
|                                    |    |     | (A)                          |           |           | (B)                           |      |             |              |
|                                    |    |     | .....ELAPSED TIME.....       |           |           | NOT IWAIT TIME(ELAPSED-IWAIT) |      |             | DISTRIBUTION |
| OCCURRENCES                        |    |     | TOTAL                        | MEAN      | MAXIMUM   | TOTAL                         | MEAN | MAXIMUM     | NUMBER       |
| SCHEDULING AND TERMINATION         |    |     |                              |           |           |                               |      |             |              |
| **REGION                           | 5  | 5   | 4146                         | 829       | 948       | 4146                          | 829  | 948         | 287A,B       |
| **REGION                           | 6  | 7   | 6028                         | 861       | 1067      | 6028                          | 861  | 1067        | 214A,B       |
| **REGION                           | 8  | 8   | 6847                         | 855       | 1098      | 6847                          | 855  | 1098        | 129A,B       |
| **REGION                           | 10 | 7   | 9664                         | 1380      | 3668      | 9664                          | 1380 | 3668        | 272A,B       |
| **REGION                           | 47 | 6   | 5482                         | 913       | 1021      | 5482                          | 913  | 1021        | 145A,B       |
| **REGION                           | 49 | 3   | 2612                         | 870       | 917       | 2612                          | 870  | 917         | 443A,B       |
| **TOTALS                           |    | 123 | 126042                       | 1024      |           | 126042                        | 1024 |             |              |
| SCHEDULE TO FIRST CALL             |    |     |                              |           |           |                               |      |             |              |
| **REGION                           | 1  | 1   | 15479797                     | 15479797  | 15479797  |                               |      |             | 555          |
| **REGION                           | 2  | 1   | 22376350                     | 22376350  | 22376350  |                               |      |             | 564          |
| **REGION                           | 3  | 1   | 15169488                     | 15169488  | 15169488  |                               |      |             | 578          |
| **REGION                           | 4  | 1   | 48146258                     | 48146258  | 48146258  |                               |      |             | 584          |
| **REGION                           | 48 | 1   | 795351                       | 795351    | 795351    |                               |      |             | 592          |
| **REGION                           | 49 | 4   | 2960425                      | 740106    | 2951746   |                               |      |             | 442          |
| **REGION                           | 50 | 1   | 15713464                     | 15713464  | 15713464  |                               |      |             | 575          |
| **TOTALS                           |    | 168 | 514286738                    | 3061230   |           |                               |      |             |              |
| ELAPSED EXECUTION                  |    |     |                              |           |           |                               |      |             |              |
| **REGION                           | 1  | 1   | 290146255                    | 290146255 | 290146255 |                               |      |             | 1            |
| **REGION                           | 2  | 1   | 252290108                    | 252290108 | 252290108 |                               |      |             | 2            |
| **REGION                           | 3  | 1   | 259496970                    | 259496970 | 259496970 |                               |      |             | 3            |
| **REGION                           | 4  | 1   | 322812716                    | 322812716 | 322812716 |                               |      |             | 4            |
| **REGION                           | 48 | 1   | 273871107                    | 273871107 | 273871107 |                               |      |             | 48           |
| **REGION                           | 49 | 4   | 271703421                    | 67925855  | 155176058 |                               |      |             | 49           |
| **REGION                           | 50 | 1   | 290379922                    | 290379922 | 290379922 |                               |      |             | 50           |
| **TOTALS                           |    | 173 | 14238540145                  | 82303700  |           |                               |      |             |              |
| DL/I CALLS                         |    |     |                              |           |           |                               |      |             |              |
|                                    |    |     |                              |           |           |                               |      | IWT/CALL(C) |              |

|          |    |       |             |         |          |             |         |          |      |          |
|----------|----|-------|-------------|---------|----------|-------------|---------|----------|------|----------|
| **REGION | 1  | 60    | 264626241   | 4410437 | 88981490 | 263813671   | 4396894 | 88970053 | 0.76 | 247A,B,C |
| **REGION | 2  | 223   | 230505269   | 1033655 | 61048758 | 227368742   | 1019590 | 61011153 | 0.73 | 237A,B,C |
| **REGION | 3  | 29    | 257704383   | 8886358 | 69000514 | 257704383   | 8886358 | 69000514 | 0.00 | 98A,B,C  |
| **REGION | 4  | 792   | 313735347   | 396130  | 52439653 | 310609035   | 392183  | 52439653 | 0.22 | 180A,B,C |
| **REGION | 49 | 592   | 262886317   | 444064  | 30202068 | 234394017   | 395935  | 30159782 | 2.46 | 177A,B,C |
| **REGION | 50 | 36    | 242591451   | 6738651 | 48651260 | 242591451   | 6738651 | 48651260 | 0.00 | 289A,B,C |
| **TOTALS |    | 18632 | 12386905286 | 664818  |          | 12024562411 | 645371  |          | 0.97 |          |

IDLE FOR INTENT

CHECKPOINT

REGION OCCUPANCY

|          |    |        |
|----------|----|--------|
| **REGION | 1  | 100.0% |
| **REGION | 2  | 100.0% |
| **REGION | 3  | 100.0% |
| **REGION | 4  | 100.0% |
| **REGION | 48 | 100.0% |
| **REGION | 49 | 100.0% |
| **REGION | 50 | 100.0% |

The following example shows a region wait report.

IMS MONITOR \*\*\*\*\*REGION IWAIT\*\*\*\*\* TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0023

| **REGION                 | 5 OCCURRENCES | TOTAL   | MEAN   | MAXIMUM | FUNCTION        | MODULE    | DISTRIBUTION NUMBER |
|--------------------------|---------------|---------|--------|---------|-----------------|-----------|---------------------|
| SCHEDULING + TERMINATION |               |         |        |         |                 |           |                     |
| SUB-TOTAL                |               |         |        |         |                 |           |                     |
| TOTAL                    |               |         |        |         |                 |           |                     |
| DL/I CALLS               |               |         |        |         |                 |           |                     |
|                          | 11            | 181816  | 16528  | 24375   | DD=IMMSTR2A     | DBH       | 117                 |
|                          | 1             | 1667    | 1667   | 1667    | PSB=BMPFPE06    | BLR-64BIT | 11                  |
|                          | 1             | 1635    | 1635   | 1635    | INT=BMPFPE06    | BLR-64BIT | 12                  |
|                          | 8             | 112831  | 14103  | 17846   | DD=IMMSTR1A     | DBH       | 118                 |
|                          | 5             | 85460   | 17092  | 33717   | DD=IMMSTR3A     | DBH       | 119                 |
|                          | 5             | 58420   | 11684  | 14643   | DD=IMINDEXA     | VBH       | 120                 |
|                          | 12            | 173866  | 14488  | 22152   | DD=PRODCNTA     | VBH       | 121                 |
|                          | 3             | 100576  | 33525  | 68373   | DD=IMMSTR2B     | DBH       | 428                 |
|                          | 1             | 17921   | 17921  | 17921   | DD=IMMSTR3B     | DBH       | 429                 |
|                          | 1             | 17195   | 17195  | 17195   | DD=IMMSTR1B     | DBH       | 430                 |
|                          | 1             | 13577   | 13577  | 13577   | DD=IMINDEXB     | VBH       | 431                 |
|                          | 3             | 49928   | 16642  | 20396   | DD=PRODCNTB     | VBH       | 432                 |
|                          | 4             | 10973   | 2743   | 2787    | DD=ITEMACTB     | DBH       | 453                 |
|                          | 2             | 37680   | 18840  | 27664   | DD=IAINDEXB     | VBH       | 454                 |
|                          | 49            | 1500067 | 30613  | 138284  | DD=INVENTRA     | DBH       | 472                 |
|                          | 23            | 345595  | 15025  | 27613   | DD=VENDORDA     | VBH       | 473                 |
|                          | 1             | 342952  | 342952 | 342952  | PI=VENDORDA...1 |           | 498                 |
|                          | 1             | 14612   | 14612  | 14612   | PI=VNSINDXA...1 |           | 499                 |
|                          | 6             | 69203   | 11533  | 19492   | DD=VNSINDXA     | VBH       | 500                 |
| TOTAL                    | 136           | 3132672 | 23034  |         |                 |           |                     |

The following example shows a programs-by-region report.

IMS MONITOR \*\*\*\*\*PROGRAMS BY REGION\*\*\*\*\* TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0069

| **REGION      | OCCURRENCES | ELAPSED EXECUTION TIME | SCHEDULING END TO FIRST CALL | DISTRIBUTION NUMBER |
|---------------|-------------|------------------------|------------------------------|---------------------|
|               |             | TOTAL MEAN MAXIMUM     | TOTAL MEAN MAXIMUM           |                     |
| 1             | 1           |                        |                              |                     |
| PROGSC6D      | 1           | 290146255              | 290146255                    | 290146255           |
| REGION TOTALS | 1           | 290146255              | 290146255                    | 290146255           |
| 2             | 2           |                        |                              |                     |
| PROGIT8C      | 1           | 252290108              | 252290108                    | 252290108           |
| REGION TOTALS | 1           | 252290108              | 252290108                    | 252290108           |
| 3             | 3           |                        |                              |                     |
| PROGTS1C      | 1           | 259496970              | 259496970                    | 259496970           |
| REGION TOTALS | 1           | 259496970              | 259496970                    | 259496970           |
| 4             | 4           |                        |                              |                     |
| PROGPS3D      | 1           | 322812716              | 322812716                    | 322812716           |
|               |             |                        |                              |                     |

|               |   |           |           |           |          |          |                  |
|---------------|---|-----------|-----------|-----------|----------|----------|------------------|
| REGION TOTALS | 1 | 322812716 | 322812716 |           | 48146258 | 48146258 |                  |
| **REGION      | 5 |           |           |           |          |          |                  |
|               | 5 |           |           |           |          |          |                  |
| PROGPS3A      | 2 | 62893103  | 31446551  | 40693590  | 5435     | 2717     | 2862 901A,B      |
| PROGTS1B      | 1 | 61794787  | 61794787  | 61794787  | 2790     | 2790     | 2790 1271A,B     |
| PROGPS3B      | 1 | 18294458  | 18294458  | 18294458  | 3104     | 3104     | 3104 1350A,B     |
| PROGIT2B      | 1 | 36095342  | 36095342  | 36095342  | 2731     | 2731     | 2731 1363A,B     |
| PROGSC2A      | 1 | 93902771  | 93902771  | 93902771  | 1667791  | 1667791  | 1667791 1401A,B  |
| REGION TOTALS | 6 | 272980461 | 45496743  |           | 1681851  | 280308   |                  |
| **REGION      | 6 |           |           |           |          |          |                  |
|               | 6 |           |           |           |          |          |                  |
| PROGIT1B      | 2 | 39000315  | 19500157  | 23703429  | 5286     | 2643     | 2801 905A,B      |
| PROGTS1B      | 1 | 34293636  | 34293636  | 34293636  | 3136     | 3136     | 3136 1207A,B     |
| PROGPS3A      | 1 | 51887767  | 51887767  | 51887767  | 2534     | 2534     | 2534 1278A,B     |
| PROGPS3B      | 2 | 67375031  | 33687515  | 40291430  | 17210570 | 8605285  | 17213287 1328A,B |
| PROGIT8A      | 1 | 69132416  | 69132416  | 69132416  | 3291     | 3291     | 3291 1359A,B     |
| PROGSC4A      | 1 | 30165017  | 30165017  | 30165017  | 2571     | 2571     | 2571 1433A,B     |
| REGION TOTALS | 8 | 291854182 | 36481772  |           | 17193752 | 2149219  |                  |
| **REGION      | 7 |           |           |           |          |          |                  |
|               | 7 |           |           |           |          |          |                  |
| PROGSC2B      | 1 | 269618583 | 269618583 | 269618583 | 5047875  | 5047875  | 5047875 909A,B   |
| REGION TOTALS | 1 | 269618583 | 269618583 |           | 5047875  | 5047875  |                  |
| **REGION      | 8 |           |           |           |          |          |                  |
|               | 8 |           |           |           |          |          |                  |
| PROGIT8A      | 1 | 5181039   | 5181039   | 5181039   | 2928     | 2928     | 2928 913A,B      |
| PROGPS3A      | 1 | 27304257  | 27304257  | 27304257  | 3350     | 3350     | 3350 1132A,B     |
| PROGSC4B      | 1 | 37286872  | 37286872  | 37286872  | 3009     | 3009     | 3009 1255A,B     |
| PROGIT2A      | 1 | 36902995  | 36902995  | 36902995  | 2850     | 2850     | 2850 1298A,B     |
| PROGIT1B      | 1 | 30407479  | 30407479  | 30407479  | 2565     | 2565     | 2565 1336A,B     |
| PROGIT1A      | 3 | 109875360 | 36625120  | 45190114  | 4279008  | 1426336  | 4272096 1357A,B  |
| PROGIT8B      | 1 | 23405220  | 23405220  | 23405220  | 2679     | 2679     | 2679 1395A,B     |
| REGION TOTALS | 9 | 270363222 | 30040358  |           | 4296389  | 477376   |                  |

## Detecting database processing intent conflicts

The IMS Monitor records the intervals when a region is in an idle state waiting to update a database owned exclusively by another already scheduled application program.

You can see the total, maximum, and average idle times in IDLE FOR INTENT following the DL/I calls. The elapsed time during the unsuccessful scheduling of a program in that region is included in the summary line times for that region.

The region can fail to be scheduled even when ownership of that database is released. The number of times processing is held up by intent failure is separately tallied under the title INTENT FAILURE SUMMARY. The report is illustrated in “Detecting database processing intent conflicts” on page 743. This report shows which PSBs are in conflict because of exclusive intent for a segment type and the database name in question.

## Examining the effects of checkpoints

The checkpoint line of the Region Summary report at the end of the region 0 summary, shows the following:

- The number of times that a system checkpoint is taken during the monitor interval
- The elapsed times
- The not-wait times

Checkpoint processing can be initiated by the control program at a specified frequency determined by the number of records placed on the system log. Other checkpoints can be caused by operator commands.

The wait time experienced during checkpoints is reported at the end of the first region summary on the Region Wait report. You can detect delays for each combination of DD name and module code. Typical entries here are for the message queue data sets and the restart data set. If a wait for storage is the cause, the entry under the FUNCTION column is STG.= followed by the identification of the pool.

**Measuring region occupancy**

Region occupancy shows the ratio of elapsed time a PST spent processing UORs to the total time of the monitor interval.

**Monitoring application program elapsed time**

The IMS Monitor can record measurements of elapsed times for each UOR. It does this during the monitored interval while other UORs are executing concurrently. Elapsed times are calculated from the start of the first DL/I (or other) call to the end of that program.

You can distinguish between time spent in application code and in DL/I processing. The event intervals are illustrated in the following figure.




Figure 91. Event intervals

Within the elapsed time for a DL/I call, the wait time to obtain segment data is recorded separately. Also, the elapsed time from schedule to first call is separately recorded. This time covers the processing in the CCTL and the database resource adapter (DRA).

The elapsed times are given in the Program Summary report. The following example shows the report. Programs are identified by their PSB name on individual lines in the report. Each line gives a summary of the activity for that



PSB during the measured interval. The total number of schedules, DL/I calls, transactions dequeued, and waits for a DL/I call for I/O are given. The report line gives calculated average times for elapsed time per schedule, processor time per schedule, schedule to first DL/I call per schedule, and elapsed time per transaction. Frequencies for calls per transaction, I/O waits per DL/I call, and transactions dequeued per schedule are also given. A TOTALS line summarizes all activity for the PSBs active during the monitored interval. (The PSB DUMMY line is a reconciliation for any incomplete scheduling caused by a region being stopped during scheduling or for a program that experiences a pseudo-abend.)

In this report, *transaction* and *schedule* can be interpreted as UOR.

| IMS MONITOR |                | ****PROGRAM SUMMARY**** |       |                |               | TRACE START 1993 130 5:55:15 |                                |                               |               | TRACE STOP 1993 130 5:59:49       |                                         |               |                            | PAGE 0075 |  |
|-------------|----------------|-------------------------|-------|----------------|---------------|------------------------------|--------------------------------|-------------------------------|---------------|-----------------------------------|-----------------------------------------|---------------|----------------------------|-----------|--|
| PSBNAME     | NO.<br>SCHEDS. | TRANS.<br>DEQ.          | CALLS | CALLS<br>/TRAN | I/O<br>IWAITS | I/O<br>IWAITS<br>/CALL       | (A)<br>TRAN.<br>DEQ./<br>/SCH. | (B)<br>CPU<br>TIME<br>/SCHED. | DISTR.<br>NO. | (A)<br>ELAPSED<br>TIME<br>/SCHED. | (B)<br>SCHED. TO<br>1ST CALL<br>/SCHED. | DISTR.<br>NO. | ELAPSED<br>TIME<br>/TRANS. |           |  |
|             |                |                         |       |                |               |                              |                                |                               |               |                                   |                                         |               |                            |           |  |
| PROGSC6D    | 1              | 13                      | 60    | 4.6            | 46            | 0.7                          | 13.0                           | 10010                         | 884A,B        | 290146255                         | 15479797                                | 886A,B        | 2231894                    |           |  |
| PROGIT8C    | 3              | 17                      | 225   | 13.2           | 166           | 0.7                          | 5.6                            | 90592                         | 888A,B        | 256617508                         | 73283259                                | 890A,B        | 45285442                   |           |  |
| PROGTS1C    | 2              | 25                      | 47    | 1.8            | 0             | 0.0                          | 12.5                           | 10010                         | 892A,B        | 239190808                         | 7586234                                 | 894A,B        | 19135264                   |           |  |
| PROGPS3D    | 1              | 23                      | 792   | 34.4           | 182           | 0.2                          | 23.0                           | 10010                         | 896A,B        | 322812716                         | 48146258                                | 898A,B        | 14035335                   |           |  |
| PROGPS3A    | 13             | 36                      | 1246  | 34.6           | 267           | 0.2                          | 2.7                            | 49782                         | 900A,B        | 32801812                          | 2228611                                 | 902A,B        | 11845098                   |           |  |
| PROGIT1B    | 11             | 21                      | 99    | 4.7            | 0             | 0.0                          | 1.9                            | 6341                          | 904A,B        | 23212388                          | 2036217                                 | 906A,B        | 12158870                   |           |  |
| PROGSC2B    | 7              | 155                     | 3068  | 19.7           | 1845          | 0.6                          | 22.1                           | 346112                        | 908A,B        | 93655514                          | 789390                                  | 910A,B        | 4229603                    |           |  |
| PROGIT8A    | 12             | 28                      | 434   | 15.5           | 293           | 0.6                          | 2.3                            | 34350                         | 912A,B        | 30196795                          | 1745815                                 | 914A,B        | 12941483                   |           |  |
| PROGPS2C    | 1              | 10                      | 179   | 17.9           | 205           | 1.1                          | 10.0                           | 10010                         | 916A,B        | 221024429                         | 53642029                                | 918A,B        | 22102442                   |           |  |
| PROGTS1B    | 8              | 20                      | 54    | 2.7            | 0             | 0.0                          | 2.5                            | 5447                          | 920A,B        | 39943245                          | 2895                                    | 922A,B        | 15977298                   |           |  |
| PROGPS3C    | 1              | 14                      | 468   | 33.4           | 117           | 0.2                          | 14.0                           | 10010                         | 924A,B        | 310644485                         | 35978027                                | 926A,B        | 22188891                   |           |  |
| PROGIT1C    | 1              | 9                       | 32    | 3.5            | 0             | 0.0                          | 9.0                            | 10010                         | 930A,B        | 304892631                         | 30226173                                | 932A,B        | 33876959                   |           |  |
| PROGSC2C    | 1              | 9                       | 160   | 17.7           | 101           | 0.6                          | 9.0                            | 10010                         | 934A,B        | 296909110                         | 22242652                                | 936A,B        | 32989901                   |           |  |
| PROGIT2B    | 8              | 21                      | 393   | 18.7           | 63            | 0.1                          | 2.6                            | 21703                         | 938A,B        | 35126671                          | 1798496                                 | 940A,B        | 13381589                   |           |  |
| PROGIT2C    | 6              | 17                      | 211   | 12.4           | 39            | 0.1                          | 2.8                            | 13312                         | 942A,B        | 288883508                         | 50698467                                | 944A,B        | 101958885                  |           |  |
| PROGTS1D    | 2              | 26                      | 50    | 1.9            | 0             | 0.0                          | 13.0                           | 10010                         | 950A,B        | 284944505                         | 10613350                                | 952A,B        | 21918808                   |           |  |
| PROGPS3B    | 8              | 22                      | 770   | 35.0           | 169           | 0.2                          | 2.7                            | 35737                         | 954A,B        | 38016279                          | 2149158                                 | 956A,B        | 13824101                   |           |  |
| PROGIT1A    | 11             | 24                      | 106   | 4.4            | 0             | 0.0                          | 2.1                            | 7925                          | 958A,B        | 30883486                          | 1935855                                 | 960A,B        | 14154931                   |           |  |
| PROGSC4A    | 9              | 163                     | 1775  | 10.8           | 5101          | 2.8                          | 18.1                           | 235921                        | 963A,B        | 62172947                          | 3011199                                 | 965A,B        | 3432862                    |           |  |
| PROGSC6C    | 1              | 10                      | 44    | 4.4            | 38            | 0.8                          | 10.0                           | 10010                         | 967A,B        | 228098334                         | 46568124                                | 969A,B        | 22809833                   |           |  |
| PROGPS2B    | 11             | 28                      | 557   | 19.8           | 604           | 1.0                          | 2.5                            | 35069                         | 971A,B        | 33309266                          | 1181831                                 | 973A,B        | 13085783                   |           |  |
| PROGIT8D    | 1              | 12                      | 175   | 14.5           | 133           | 0.7                          | 12.0                           | 10010                         | 975A,B        | 253392289                         | 21274169                                | 977A,B        | 21116024                   |           |  |
| PROGSC4C    | 1              | 10                      | 98    | 9.8            | 349           | 3.5                          | 10.0                           | 10010                         | 979A,B        | 248736332                         | 25930126                                | 981A,B        | 34873633                   |           |  |
| PROGSC6A    | 7              | 157                     | 789   | 5.0            | 457           | 0.5                          | 22.4                           | 11703                         | 983A,B        | 73936039                          | 115979                                  | 985A,B        | 3296511                    |           |  |
| PROGIT2A    | 7              | 22                      | 430   | 19.5           | 71            | 0.1                          | 3.1                            | 28529                         | 987A,B        | 37905001                          | 2982                                    | 989A,B        | 12060682                   |           |  |
| PROGSC2D    | 1              | 15                      | 280   | 18.6           | 180           | 0.6                          | 15.0                           | 10010                         | 991A,B        | 316194222                         | 41527764                                | 993A,B        | 21079614                   |           |  |
| PROGPS2A    | 6              | 25                      | 490   | 19.6           | 548           | 1.1                          | 4.1                            | 43177                         | 995A,B        | 58277945                          | 2467506                                 | 997A,B        | 13986707                   |           |  |
| PROGSC2A    | 5              | 121                     | 2363  | 19.5           | 1420          | 0.6                          | 24.2                           | 276187                        | 1001A,B       | 88906184                          | 6022954                                 | 1003A,B       | 3673809                    |           |  |
| PROGIT2D    | 1              | 20                      | 361   | 18.0           | 62            | 0.1                          | 20.0                           | 10010                         | 1005A,B       | 386092737                         | 111426279                               | 1007A,B       | 19304636                   |           |  |
| PROGSC4B    | 10             | 131                     | 1421  | 10.8           | 4115          | 2.8                          | 13.1                           | 617016                        | 1011A,B       | 53826667                          | 2632409                                 | 1013A,B       | 4108905                    |           |  |
| PROGSC4D    | 1              | 19                      | 197   | 10.3           | 668           | 3.3                          | 19.0                           | 10010                         | 1020A,B       | 227999124                         | 46667334                                | 1022A,B       | 11999953                   |           |  |
| PROGPS2D    | 1              | 13                      | 240   | 18.4           | 291           | 1.2                          | 13.0                           | 10010                         | 1025A,B       | 327602445                         | 52935987                                | 1027A,B       | 25200188                   |           |  |
| PROGSC6B    | 5              | 140                     | 694   | 4.9            | 395           | 0.5                          | 28.0                           | 16884                         | 1032A,B       | 78994223                          | 3290769                                 | 1034A,B       | 2821222                    |           |  |
| PROGIT1D    | 1              | 10                      | 36    | 3.6            | 0             | 0.0                          | 10.0                           | 10010                         | 1041A,B       | 290379922                         | 15713464                                | 1043A,B       | 29037992                   |           |  |
| PROGIT8B    | 8              | 17                      | 288   | 16.9           | 190           | 0.6                          | 2.1                            | 33436                         | 1259A,B       | 35223857                          | 2902                                    | 1261A,B       | 16575932                   |           |  |
| **TOTALS    | 173            | 1403                    | 18632 | 13.2           | 18115         | 0.9                          | 8.1                            | 90328                         |               | 82303700                          | 2972755                                 |               | 10148638                   |           |  |

To examine the detail of the call processing for each program (itemized by type of call and summarized for the monitor interval), you can use the Call Summary report. An extract from the multipage output is given in the following example. The calls using an I/O PCB are given first and sub-totaled. Then, the total calls, of each type, against each database PCB and each external subsystem are listed. The PSB TOTAL line marks the end of data for each program.

| IMS MONITOR |          | ****CALL SUMMARY**** |         |         |           | TRACE START 1993 130 5:55:15 |        |              |                | TRACE STOP 1993 130 5:59:49 |           |               |                | PAGE 0186       |  |
|-------------|----------|----------------------|---------|---------|-----------|------------------------------|--------|--------------|----------------|-----------------------------|-----------|---------------|----------------|-----------------|--|
| PSB NAME    | PCB NAME | CALL FUNC            | LEV NO. | SEGMENT | STAT CODE | CALLS                        | IWAITS | IWAITS/ CALL | ..ELAPSED MEAN | TIME... MAXIMUM             | .NOT MEAN | IWAIT MAXIMUM | TIME.. MAXIMUM | DISTRIB. NUMBER |  |
|             |          |                      |         |         |           |                              |        |              |                |                             |           |               |                |                 |  |
| PROGSC6B    | I/O PCB  | ISRT ( )             |         |         |           | 138                          | 0      | 0.00         | 372            | 1240                        | 372       | 1240          | 598A,B,C       |                 |  |
|             |          | GU ( )               |         |         |           | 134                          | 133    | 0.99         | 2600917        | 20974615                    | 2587532   | 20962866      | 602A,B,C       |                 |  |
|             |          | (GU) ( )             |         |         |           | 3                            | 0      | 0.00         | 15             | 16                          | 15        | 16            | 716A,B,C       |                 |  |
|             |          | ASRT ( )             |         |         |           | 3                            | 0      | 0.00         | 330            | 333                         | 330       | 333           | 869A,B,C       |                 |  |
|             |          | GU ( )               |         |         | QC        | 2                            | 1      | 0.50         | 17639806       | 21219588                    | 17634776  | 21209529      | 870A,B,C       |                 |  |
|             |          | I/O PCB SUBTOTAL     |         |         |           |                              |        |              |                |                             |           |               |                |                 |  |
|             |          |                      |         |         |           | 280                          | 134    | 0.47         | 1370910        |                             | 1364469   |               |                |                 |  |
|             | INVENTRB | DLET (03)IN060SUP    |         |         |           | 138                          | 0      | 0.00         | 813            | 1289                        | 813       | 1289          | 599A,B,C       |                 |  |
|             |          | GNP (03)IN060SUP     |         |         |           | 138                          | 7      | 0.05         | 2112           | 112589                      | 1047      | 112589        | 600A,B,C       |                 |  |
|             |          | GU (01)IN010PAR      |         |         |           | 138                          | 254    | 1.84         | 29511          | 75356                       | 1195      | 19229         | 601A,B,C       |                 |  |
|             |          | DL/I PCB SUBTOTAL    |         |         |           |                              |        |              |                |                             |           |               |                |                 |  |
|             |          |                      |         |         |           | 414                          | 261    | 0.63         | 10812          |                             | 1018      |               |                |                 |  |



|                  |                   |      |      |      |          |          |          |          |          |
|------------------|-------------------|------|------|------|----------|----------|----------|----------|----------|
| PSB TOTAL        |                   | 694  | 395  | 0.56 | 559555   |          | 551114   |          |          |
| PROGSC2A I/O PCB | TSRT ( )          | 118  | 0    | 0.00 | 381      | 1496     | 381      | 1496     | 603A,B,C |
|                  | GU ( )            | 114  | 284  | 2.49 | 3304809  | 21784513 | 3164423  | 21664181 | 632A,B,C |
|                  | (GU) ( )          | 2    | 0    | 0.00 | 17       | 18       | 17       | 18       | 781A,B,C |
|                  | ASRT ( )          | 3    | 0    | 0.00 | 367      | 444      | 367      | 444      | 871A,B,C |
|                  | GU ( )            | 2    | 5    | 2.50 | 19931897 | 20045206 | 19799530 | 19925277 | 872A,B,C |
|                  | I/O PCB SUBTOTAL  |      |      |      |          |          |          |          |          |
|                  |                   | 239  | 289  | 1.20 | 1743339  |          | 1675270  |          |          |
| LOGVENDA         | REPL (03)IN040SLQ | 118  | 0    | 0.00 | 268      | 804      | 268      | 804      | 604A,B,C |
|                  | GNP (03)IN040SLQ  | 118  | 5    | 0.04 | 899      | 16995    | 218      | 305      | 605A,B,C |
|                  | REPL (02)VN030PAR | 826  | 0    | 0.00 | 805      | 1578     | 805      | 1578     | 606A,B,C |
|                  | GNP (02)VN030PAR  | 826  | 873  | 1.05 | 19321    | 94521    | 456      | 1363     | 607A,B,C |
|                  | REPL (01)VN020REO | 118  | 58   | 0.49 | 8879     | 48076    | 832      | 1682     | 623A,B,C |
|                  | GU (01)VN020REO   | 118  | 195  | 1.65 | 31688    | 360775   | 1300     | 1746     | 625A,B,C |
|                  | DL/I PCB SUBTOTAL |      |      |      |          |          |          |          |          |
|                  |                   | 2124 | 1131 | 0.53 | 10145    |          | 636      |          |          |
| PSB TOTAL        |                   | 2363 | 1420 | 0.60 | 185445   |          | 170013   |          |          |
| PROGSC2D I/O PCB | TSRT ( )          | 14   | 0    | 0.00 | 377      | 621      | 377      | 621      | 608A,B,C |
|                  | GU ( )            | 14   | 36   | 2.57 | 22360408 | 52048566 | 22221852 | 51901313 | 634A,B,C |
|                  | I/O PCB SUBTOTAL  |      |      |      |          |          |          |          |          |
|                  |                   | 28   | 36   | 1.28 | 11180393 |          | 11111115 |          |          |
| LOGVENDD         | REPL (03)IN040SLQ | 14   | 0    | 0.00 | 263      | 328      | 263      | 328      | 609A,B,C |
|                  | GNP (03)IN040SLQ  | 14   | 1    | 0.07 | 1407     | 16889    | 223      | 307      | 610A,B,C |
|                  | REPL (02)VN030PAR | 98   | 0    | 0.00 | 820      | 1015     | 820      | 1015     | 611A,B,C |

## Monitoring I/O for application program DL/I calls

The IMS Monitor report shows the total number of I/O occurrences and the total time the occurrences took for each UOR during a monitored interval.

The Program I/O report gives these two totals for all PSBs active during the monitored interval and includes the detailed breakdown of the I/O wait time as it was incurred by each PCB used by the program.

The report shows any contention experienced during application program processing. Each type of conflict and the number of times it occurred are recorded for each I/O PCB or database PCB. The report shows the total wait time, the highest wait experienced, and the average time. Subtotals are given for each PCB under a PSB, and for all PCBs under each PSB.

The DDN/FUNC column list the data set by DD name. The MODULE column uses a code to indicate the source of the contention. The types of conflicts and codes are shown as follows. Any codes that appear apply to IMS only.

- **Scheduling**

- Code Conflict**

- BLR** Load/read from ACBLIB

- SMN** Virtual storage management

- **Database access**

- Code Conflict**

- DBH** OSAM I/O

- DLE** DL/I functions

- VBH** VSAM interface

- (Physical segment code)**

- Program isolation

The I/O waits for the calls to the I/O PCB are grouped as the first entries for a PSB. For DL/I calls, the data set for which the I/O took place is indicated under the DDN/FUNC heading, and the module code tells you what type of conflict caused the wait.

Contention for the same physical segment in a database causes a wait on behalf of program isolation. This is shown in the DDN/FUNC column, on the PCB line, by the entry PIdmb, where dmb is the DMB of the physical data set. The MODULE column identifies the segment type using the physical segment code assigned by DBD generation.

When an application is accessing a database using VSAM as the access method, DL/I calls do not generally result in an I/O wait. A MODULE column entry of VBH indicates that interface to VSAM occurred and there was an I/O wait.

A seemingly unrelated entry can occur under the DDN/FUNC column for a database PCB. An example is a retrieval call to a database (DB-A) that causes a buffer to be purged in order to make room for that retrieved data. If the buffer contents included data belonging to another database (DB-B), the I/O entry in the report shows the DD name for DB-B as being in conflict for PCB access to DB-A.

## Transaction Queuing report

The Transaction Queuing Report in the following example lists the transactions for DBCTL. Each transaction name is an 8-byte transaction ID specified by the CCTL on the schedule request or the CCTL ID.

A transaction ID from CICS, when used as the transaction manager, is composed of a 4-byte CICS transaction name plus a 4-byte CICS identifier. If the CCTL does not specify the transaction ID, DBCTL takes the CCTL region ID obtained at connection time as the default. In this report for DBCTL, the transaction NUMBER DEQUEUED is the number of schedules, and the ON QUEUE WHEN SCHEDULED is always zero, because the IMS message queues are not involved.

| IMS MONITOR |                 | ****TRANSACTION QUEUING**** |            |      | TRACE START 1993 130 5:55:15 |       | TRACE STOP 1993 130 5:59:49 |  | PAGE 0181 |
|-------------|-----------------|-----------------------------|------------|------|------------------------------|-------|-----------------------------|--|-----------|
| TRANSACTION | NUMBER DEQUEUED | NUMBER SCHEDS.              | (B)        |      | SCHEDULED.....               | (A)   | DISTRIBUTION                |  |           |
|             |                 |                             | ..ON QUEUE | WHEN |                              |       |                             |  |           |
|             |                 |                             | MINIMUM    | MEAN | MAXIMUM                      | MEAN  | NUMBER                      |  |           |
| SC6X        | 13              | 1                           | 0          | 0.00 | 0                            | 13.00 | 883A,B                      |  |           |
| IT8W        | 17              | 3                           | 0          | 0.00 | 0                            | 5.66  | 887A,B                      |  |           |
| TS1Z        | 16              | 1                           | 0          | 0.00 | 0                            | 16.00 | 891A,B                      |  |           |
| PS3X        | 23              | 1                           | 0          | 0.00 | 0                            | 23.00 | 895A,B                      |  |           |
| PS3Y        | 17              | 7                           | 0          | 0.00 | 0                            | 2.42  | 899A,B                      |  |           |
| IT1V        | 11              | 6                           | 0          | 0.00 | 0                            | 1.83  | 903A,B                      |  |           |
| SC2Z        | 143             | 2                           | 0          | 0.00 | 0                            | 71.50 | 907A,B                      |  |           |
| IT8U        | 12              | 7                           | 0          | 0.00 | 0                            | 1.71  | 911A,B                      |  |           |
| PS2W        | 10              | 1                           | 0          | 0.00 | 0                            | 10.00 | 915A,B                      |  |           |
| TS1U        | 12              | 4                           | 0          | 0.00 | 0                            | 3.00  | 919A,B                      |  |           |
| PS3W        | 14              | 1                           | 0          | 0.00 | 0                            | 14.00 | 923A,B                      |  |           |
| IT8Y        | 16              | 5                           | 0          | 0.00 | 0                            | 3.20  | 927A,B                      |  |           |
| IT1W        | 9               | 1                           | 0          | 0.00 | 0                            | 9.00  | 929A,B                      |  |           |
| SC2W        | 9               | 1                           | 0          | 0.00 | 0                            | 9.00  | 933A,B                      |  |           |
| IT2V        | 13              | 5                           | 0          | 0.00 | 0                            | 2.60  | 937A,B                      |  |           |
| IT2W        | 17              | 6                           | 0          | 0.00 | 0                            | 2.83  | 941A,B                      |  |           |
| TS1V        | 9               | 1                           | 0          | 0.00 | 0                            | 9.00  | 945A,B                      |  |           |
| SC2V        | 12              | 5                           | 0          | 0.00 | 0                            | 2.40  | 947A,B                      |  |           |
| TS1W        | 11              | 1                           | 0          | 0.00 | 0                            | 11.00 | 949A,B                      |  |           |
| PS3V        | 13              | 3                           | 0          | 0.00 | 0                            | 4.33  | 953A,B                      |  |           |
| IT1U        | 9               | 6                           | 0          | 0.00 | 0                            | 1.50  | 957A,B                      |  |           |
| SC4U        | 11              | 5                           | 0          | 0.00 | 0                            | 2.20  | 962A,B                      |  |           |
| SC6W        | 10              | 1                           | 0          | 0.00 | 0                            | 10.00 | 966A,B                      |  |           |
| PS2V        | 8               | 6                           | 0          | 0.00 | 0                            | 1.33  | 970A,B                      |  |           |
| IT8X        | 12              | 1                           | 0          | 0.00 | 0                            | 12.00 | 974A,B                      |  |           |
| SC4W        | 10              | 1                           | 0          | 0.00 | 0                            | 10.00 | 978A,B                      |  |           |
| SC6U        | 14              | 6                           | 0          | 0.00 | 0                            | 2.33  | 982A,B                      |  |           |
| IT2Y        | 9               | 3                           | 0          | 0.00 | 0                            | 3.00  | 986A,B                      |  |           |
| SC2X        | 15              | 1                           | 0          | 0.00 | 0                            | 15.00 | 990A,B                      |  |           |
| PS2Y        | 17              | 2                           | 0          | 0.00 | 0                            | 8.50  | 994A,B                      |  |           |
| SC4Y        | 152             | 4                           | 0          | 0.50 | 1                            | 38.00 | 998A,B                      |  |           |
| SC2Y        | 106             | 2                           | 0          | 0.00 | 0                            | 53.00 | 1000A,B                     |  |           |
| IT2X        | 20              | 1                           | 0          | 0.00 | 0                            | 20.00 | 1004A,B                     |  |           |
| SC2U        | 15              | 3                           | 0          | 0.00 | 0                            | 5.00  | 1008A,B                     |  |           |

|      |     |   |   |      |   |        |         |
|------|-----|---|---|------|---|--------|---------|
| SC4Z | 123 | 5 | 0 | 0.60 | 1 | 24.60  | 1010A,B |
| TS1X | 15  | 1 | 0 | 0.00 | 0 | 15.00  | 1015A,B |
| SC4X | 19  | 1 | 0 | 0.00 | 0 | 19.00  | 1019A,B |
| PS2X | 13  | 1 | 0 | 0.00 | 0 | 13.00  | 1024A,B |
| PS2Z | 20  | 5 | 0 | 0.00 | 0 | 4.00   | 1028A,B |
| SC6Z | 130 | 1 | 0 | 0.00 | 0 | 130.00 | 1031A,B |
| SC6V | 10  | 4 | 0 | 0.00 | 0 | 2.50   | 1035A,B |
| SC6Y | 143 | 1 | 0 | 0.00 | 0 | 143.00 | 1037A,B |
| IT1X | 10  | 1 | 0 | 0.00 | 0 | 10.00  | 1040A,B |
| PS3U | 19  | 6 | 0 | 0.00 | 0 | 3.16   | 1131A,B |
| IT2U | 13  | 4 | 0 | 0.00 | 0 | 3.25   | 1146A,B |

## Monitoring database buffers

One of the key resources in an online system is the database buffer pool. The efficiency of DL/I call service depends on the presence of the required database logical record in the buffer, so that segment retrieval does not require additional I/O.

This is especially true for HOLD calls with intervening database calls prior to a replace call.

Use the QUERY POOL command to view usage statistics for Fast Path buffer pools.

You can assess the general efficiency of the full-function pool management using the Database Buffer Pool report shown in the following example. The event counts on this report are not specific to a particular database or program but represent the pressure for use of the database pool.

**Related reading:** Refer to “Database-Buffer-Pool report” on page 720 for more information on the Database Buffer Pool reports.

### DATA BASE BUFFER POOL

|                                                        | FIX PREFIX/BUFFERS           |                       | Y/Y        |
|--------------------------------------------------------|------------------------------|-----------------------|------------|
|                                                        | SUBPOOL ID                   |                       | 004K       |
|                                                        | SUBPOOL BUFFER SIZE          |                       | 4096       |
|                                                        | TOTAL BUFFERS IN SUBPOOL     |                       | 1000       |
|                                                        | 16:09:59<br>START TRACE      | 16:25:10<br>END TRACE | DIFFERENCE |
| NUMBER OF LOCATE-TYPE CALLS                            | 407636                       | 4296793               | 3889157    |
| NUMBER OF REQUESTS TO CREATE NEW BLOCKS                | 1                            | 7                     | 6          |
| NUMBER OF BUFFER ALTER CALLS                           | 75006                        | 819359                | 744353     |
| NUMBER OF PURGE CALLS                                  | 9137                         | 93881                 | 84744      |
| NUMBER OF LOCATE-TYPE CALLS, DATA ALREADY IN OSAM POOL | 313896                       | 3317264               | 3003368    |
| NUMBER OF BUFFERS SEARCHED BY ALL LOCATE-TYPE CALLS    | 453364                       | 4779327               | 4325963    |
| NUMBER OF READ I/O REQUESTS                            | 86881                        | 904487                | 817606     |
| NUMBER OF SINGLE BLOCK WRITES BY BUFFER STEAL ROUTINE  | 0                            | 0                     | 0          |
| NUMBER OF BLOCKS WRITTEN BY PURGE                      | 32629                        | 360434                | 327805     |
| NUMBER OF LOCATE CALLS WAITED DUE TO BUSY ID           | 281                          | 3173                  | 2892       |
| NUMBER OF LOCATE CALLS WAITED DUE TO BUFFER BUSY WRT   | 6                            | 180                   | 174        |
| NUMBER OF LOCATE CALLS WAITED DUE TO BUFFER BUSY READ  | 0                            | 0                     | 0          |
| NUMBER OF BUFFER STEAL/PURGE WAITED FOR OWNERSHIP RLSE | 43                           | 483                   | 440        |
| NUMBER OF BUFFER STEAL REQUESTS WAITED FOR BUFFERS     | 0                            | 0                     | 0          |
| TOTAL NUMBER OF I/O ERRORS FOR THIS SUBPOOL            | 0                            | 0                     | 0          |
| NUMBER OF BUFFERS LOCKED DUE TO WRITE ERRORS           | 0                            | 0                     | 0          |
| QUOTIENT : TOTAL NUMBER OF OSAM READS + OSAM WRITES =  | 7.02                         |                       |            |
|                                                        | TOTAL NUMBER OF TRANSACTIONS |                       |            |

If any of your databases use VSAM as access method, the IMS Monitor produces a series of reports headed VSAM BUFFER POOL, one for each subpool. The following example shows one of these reports.

\*\*\*I M S M O N I T O R\*\*\* BUFFER POOL STATISTICS

V S A M B U F F E R P O O L

|                           |       |
|---------------------------|-------|
| FIX INDEX/BLOCK/DATA      | N/Y/N |
| SHARED RESOURCE POOL ID   | VPL1  |
| SHARED RESOURCE POOL TYPE | D     |
| SUBPOOL ID                | 2     |
| SUBPOOL BUFFER SIZE       | 4096  |
| NUMBER HIPERSPACE BUFFERS | 0     |
| TOTAL BUFFERS IN SUBPOOL  | 4     |

|                                                         | 16:09:59<br>START TRACE | 16:25:10<br>END TRACE | DIFFERENCE |
|---------------------------------------------------------|-------------------------|-----------------------|------------|
| NUMBER OF RETRIEVE BY RBA CALLS RECEIVED BY BUF HNDLR   | 432                     | 6029                  | 5597       |
| NUMBER OF RETRIEVE BY KEY CALLS                         | 40857                   | 443840                | 402983     |
| NUMBER OF LOGICAL RECORDS INSERTED INTO ESDS            | 414                     | 6011                  | 5597       |
| NUMBER OF LOGICAL RECORDS INSERTED INTO KSDS            | 2132                    | 25266                 | 23134      |
| NUMBER OF LOGICAL RECORDS ALTERED IN THIS SUBPOOL       | 0                       | 0                     | 0          |
| NUMBER OF TIMES BACKGROUND WRITE FUNCTION INVOKED       | 0                       | 0                     | 0          |
| NUMBER OF SYNCHRONIZATION CALLS RECEIVED                | 6494                    | 70963                 | 64469      |
| NUMBER OF WRITE ERROR BUFFERS CURRENTLY IN THE SUBPOOL  | 0                       | 0                     | 0          |
| LARGEST NUMBER OF WRITE ERRORS IN THE SUBPOOL           | 0                       | 0                     | 0          |
| NUMBER OF VSAM GET CALLS ISSUED                         | 44249                   | 487181                | 442932     |
| NUMBER OF VSAM SCHBFR CALLS ISSUED                      | 0                       | 0                     | 0          |
| NUMBER OF TIMES CTRL INTERVAL REQUESTED ALREADY IN POOL | 11886                   | 129668                | 117782     |
| NUMBER OF CRTL INTERVALS READ FROM EXTERNAL STORAGE     | 32842                   | 363635                | 330793     |
| NUMBER OF VSAM WRITES INITIATED BY IMS                  | 2370                    | 29208                 | 26838      |
| NUMBER OF VSAM WRITES TO MAKE SPACE IN THE POOL         | 0                       | 0                     | 0          |
| NUMBER OF VSAM READS FROM HIPERSPACE BUFFERS            | 0                       | 0                     | 0          |
| NUMBER OF VSAM WRITES FROM HIPERSPACE BUFFERS           | 0                       | 0                     | 0          |
| NUMBER OF FAILED VSAM READS FROM HIPERSPACE BUFFERS     | 0                       | 0                     | 0          |
| NUMBER OF FAILED VSAM WRITES FROM HIPERSPACE BUFFERS    | 0                       | 0                     | 0          |
| QUOTIENT : TOTAL NUMBER OF VSAM READS + VSAM WRITES =   | 2.19                    |                       |            |
| TOTAL NUMBER OF TRANSACTIONS                            |                         |                       |            |

## IMS internal resource usage

There are several summary reports that you can use to examine the level of internal contention for resources.

The following list provides a brief description of these reports.

### Pool space failure report

The Pool Space Failure Summary report gives the number of times (in each region) a given amount of storage was unavailable. It shows the number of bytes, the identification of the pool, and the number of occurrences when storage was unavailable. You can use this summary to determine if you need to increase the buffer pool allocation, either by a system definition change or by overriding the number of buffers in the EXEC statements in the JCL.

The format of the report is shown in the following example.

#### POOL SPACE FAILURE SUMMARY

| POOL ID | BYTES REQ. | OCCURRENCES |
|---------|------------|-------------|
| DLMP    | 8888       | 1           |
| DLDP    | 7777       | 1           |
| TOTAL   |            | 2           |

### Programs experiencing deadlock report

The Deadlock Event Summary report records each time a pair of programs reaches a deadlock over ownership of a segment in a given database data set. Each line in the report shows the two PSBs involved and indicates which is given processing right-of-way (REQ-ING PSB) and which has to

reprocess after dynamic backout has occurred (LOSING PSB). The deadlock event summary report is illustrated in the following example.

#### DEADLOCK EVENT SUMMARY

|             |            |          |             |
|-------------|------------|----------|-------------|
| REQ-ING PSB | LOSING PSB | DMBNAME  | OCCURRENCES |
| PSBNAME1    | TPPSBRE3   | DBASEBAL | 1           |
| TOTAL       |            |          | 1           |

### IMS latch conflict report

The basic serialization of the task processing in IMS is controlled by ownership of an IMS latch. When different programs are executing, they compete for the ownership. If they wait for the resource, the one possessing the latch has to post the other ITASK waiting for it. Use the Latch Conflict Statistics report to judge the level of contention for a resource.

The different types of latches and the counters that exhibit the level of contention are given in the Latch Conflict Statistics report. The following example shows a sample latch conflict statistics report. The entries are organized according to the latch names.

When a system checkpoint is taken during the time the monitor is active, latch conflict statistics are reset to zero, thus corrupting the values presented in this report. If this situation exists, the following message will be inserted at the top of the report:

```

**** A CHECKPOINT OCCURRED DURING MONITOR RUN ****
**** LATCH CONFLICT STATISTICS ARE INVALID ****
**** SEE UTILITIES REFERENCE MANUAL ****

```

However, if the master terminal operator issues the /CHECKPOINT command with the STATISTICS keyword parameter, latch conflict statistics are reset to zero, but the IMS monitor is not notified. Therefore, DFSUTR20 cannot detect that the statistics have been corrupted and will not issue this message.

**Recommendation:** Do not issue statistics checkpoints while the monitor is running.

```

IMS MONITOR ** GENERAL REPORTS **          TRACE START 1993 209...

```

#### LATCH CONFLICT STATISTICS

| LATCH<br>NAMES | COUNT<br>FIELD | AT<br>START | AT<br>END | DIFF. |
|----------------|----------------|-------------|-----------|-------|
| LOGL           | CONTENTIONS    | 0           | 0         | 0     |
| SMGT           | CONTENTIONS    | 0           | 0         | 0     |
| XCNQ           | CONTENTIONS    | 0           | 0         | 0     |
| ACTL           | CONTENTIONS    | 0           | 0         | 0     |
| CBTS           | CONTENTIONS    | 0           | 0         | 0     |
| DBLK           | CONTENTIONS    | 0           | 0         | 0     |

## Using frequency distributions from IMS Monitor output

The reports derived from the IMS Monitor data records contain many summary lines where the mean time is given.

If you are interested in the distribution of those timed events, rather than just average and maximum times, you can request the Report Print utility to individually record the events in a frequency distribution across a range of intervals. Some distributions are not time dependent, such as those for transaction queue loads or transmitted block sizes.

- The following table shows the report distributions sorted by Region Summary.

*Table 70. Report distributions by region summary*

| Report name                   | ID  | Description   |
|-------------------------------|-----|---------------|
| Scheduling and Termination    | D1  | Elapsed time  |
|                               | D2  | Not wait time |
| Schedule end to 1st DL/I call | D3  | N/A           |
| Elapsed execution time        | D4  | N/A           |
| DL/I calls                    | D5  | Elapsed time  |
|                               | D6  | Not wait time |
| External Subsystem calls      | D43 | Elapsed time  |
| Waits per DL/I call           | D7  | N/A           |
| Idle for intent               | D8  | N/A           |
| Checkpoint                    | D20 | Elapsed time  |
|                               | D21 | Not wait time |

- The following table shows the report distributions sorted by Programs Region.

*Table 71. Report distributions by programs region*

| Report name                   | ID  | Description |
|-------------------------------|-----|-------------|
| Elapsed execution time        | D30 | N/A         |
| Schedule and to 1st DL/I call | D31 | N/A         |

- The following table shows the report distributions by Program Summary.

*Table 72. Report distributions by program summary*

| Report name                        | ID  | Description |
|------------------------------------|-----|-------------|
| Processor time per schedule        | D15 | N/A         |
| Transactions dequeued per schedule | D14 | N/A         |
| Elapsed time per schedule          | D9  | N/A         |
| Schedule end to 1st DL/I call      | D10 | N/A         |

- The following table shows the report distributions sorted by Call Summary.

*Table 73. Report distributions by call summary*

| Report name                           | ID  | Description  |
|---------------------------------------|-----|--------------|
| PSB waits per DL/I call               | D13 | N/A          |
| PSB waits per external subsystem call | D44 | N/A          |
| PSB elapsed time per call             | D11 | N/A          |
| PSB not wait time per call            | D12 | N/A          |
| PSB external subsystem calls          | D45 | N/A          |
|                                       | N/A | Elapsed time |

- The following table lists some distributions derived from buffer pool statistics for wait times.

*Table 74. Wait time distributions*

| Function         | ID  | Module key |
|------------------|-----|------------|
| Storage          | D22 | SMN        |
| OSAM I/O         | D23 | DBH        |
| VSAM I/O         | D24 | VBH        |
| Block loader I/O | D27 | BLR        |
| HSAM I/O         | D34 | DIE        |
| PI enqueue       | D40 | None       |

## How to get a frequency distribution output

To request the IMS Monitor Report Print utility to gather distribution data, you must include a DIS input control statement. This causes all report items with an entry under a column headed MEAN to have a corresponding frequency distribution as part of the Distribution Appendix. Also, each report line includes an identifying reference number under the column headed DISTRIBUTION NUMBER so that you can locate the distribution data in the appendix, flagged by that same number.

The following tables show the major IMS Monitor reports and the type of frequency distributions generated for each report. Each type results in several distributions, depending on how many entries are in each section of the report. For each type of frequency distribution, the data is cumulated in suitable intervals or ranges. The set of ranges used for each type is given an identifier, shown in the ID column.

## How frequency distribution ranges are defined

A set of ten intervals is defined for each summary line and the occurrences falling in each interval are cumulated. The interval ranges are preset with default end points. For example, the end points, for DL/I call elapsed time are: 0, 1000, 2000, 4000, 8000, 16000, 32000, 64000, 128000, 256000, INF (all times are in milliseconds). The default end points are chosen so that they are suitable to the event. The lower limit of the first interval always defaults to zero, and the upper limit of the tenth interval is infinity (INF).

Although several types of distribution can use the same set of end points, each type is assigned a distribution identifier. You can use this to redefine the end points. To override the default end points include an input control statement to the Report Print utility. The statement specifies the type of distribution identifier and gives the desired end point values. For example, the DL/I call elapsed time end points could be respecified by:

```
D5 0,500,1000,1500,2000,4000,,,100000,500000
```

The values of the unspecified end points remain at their default values of 32000 and 64000 as does the last (INF).

The following example, which is a sample page taken from a Distribution Appendix, shows how individual distributions are numbered and how ranges vary with the type of distribution. The lines are arranged in pairs, with the second one recording the cumulated counts.

```

IMS MONITOR   ****DISTRIBUTION APPENDIX****   TRACE START 1993 130   5:55:15   TRACE STOP 1993 130   5:59:49   PAGE 0200
# 1.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 2.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 3.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 4.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 5.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 6.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 7.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 8.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 9.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 10.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 11.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 12.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 13.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 14.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 15.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
# 16.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF

```

## Default values of distribution definitions

Using an identifier provided in the frequency distribution tables and the Wait Time Distributions table, you can determine the default end points for the distribution by locating it in the following list:

**D1, D2, D5, D6, D9, D10, D11, D12, D15 D18, D19, D20, D21, D22, D25, D27 D28, D29, D30, D31, D43, and D45**

0, 1000, 2000, 4000, 8000, 16000, 32000, 64000, 128000, 256000, INF

**D3** 0, 50000, 100000, 150000, 200000, 250000, 300000, 350000, 400000, 450000, INF

**D4** 0, 200000, 400000, 600000, 800000, 1000000, 1200000, 1400000, 1600000, 1800000, INF

**D7, D13, and D44**

0, 0, 1, 2, 3, 4, 5, 6, 7, 8, INF

**D8** 0, 100000, 200000, 300000, 400000, 500000, 600000, 700000, 800000, 900000, INF

**D14, D16, D17**

0, 1, 2, 3, 4, 5, 10, 15, 30, 90, INF

**D23, D24, D26, D32, D40, D42**

0, 2000, 8000, 24000, 50000, 100000, 150000, 200000, 250000, 300000, INF

**D33, D34, D35**

0, 2000, 4000, 8000, 16000, 32000, 64000, 96000, 128000, 160000, INF

**D36, D37**

0, 10, 20, 40, 80, 100, 200, 400, 800, 1000, INF



**D38** 0, 1000, 10000, 100000, 200000, 500000, 800000, 1000000, 1500000, 2000000,  
INF

**D39** 0, 1000, 5000, 10000, 50000, 100000, 500000, 1000000, 5000000, 10000000, INF

## **Interpreting Distribution Appendix output**

You can use the detailed output in the Distribution Appendix when you suspect an unusual combination of events was reported in a report summary line.

Usually, the average and maximum times or counts are sufficient to highlight a resource usage problem. However, if you suspect the mean value to be masking an unusual distribution you can draw on the detail contained in the IMS Monitor output records.

---

## Chapter 50. IMS Monitor reports for DCCTL

DCCTL is a transaction management subsystem that has no database components. With the external subsystem (ESS) attach facility, it provides the transaction management capability for non-IMS database subsystems.

This topic describes:

- The events that the IMS Monitor collects
- The content of the reports produced by the IMS Monitor Report Print Program in a DCCTL environment

DCCTL does not change the format or usage of the IMS Monitor reports. There are reports, and fields within reports, that contain information specific to databases, and these are not applicable to the DCCTL environment. Reports that do not apply to DCCTL appear as a heading without data, or are not produced. The reports that do not apply to DCCTL include:

- Database Buffer Pool report
- VSAM Buffer Pool (DB) report
- Call Summary (DB) report
- Program I/O (DB) report

For a detailed look at the events, system activities, and use of storage areas (buffer pool or data set) for which timings are gathered by the IMS Monitor, see “Transaction flow and IMS Monitor events” on page 733.

---

### IMS Monitor trace event intervals

The IMS Monitor trace interval is defined by the master terminal operator's use of the /TRACE command between the start and stop command entries. The online IMS events are recorded in IMS Monitor records placed in the IMSMON data set.

The event timings are related to dependent region activity. The following figure shows the boundaries of the timed event intervals.

The Monitor trace interval includes the following intervals:

- Scheduling and Termination
  - Block loader busy
  - Schedule failures (PSB busy and space failure)
  - Sched/Term elapsed
    - NOT-WAIT
    - ACBLIB waits
- Region occupancy (which overlaps with all of Sched/Term elapsed)
  - Schedule to first call
  - Elapsed execution

The NOT-WAIT time for a region is the elapsed time not accounted for by wait time. Any delay coming from either paging or the processor being dispatched for a higher priority task results in an increase in the NOT-WAIT times.




Figure 92. IMS Monitor trace event intervals

## Overview of IMS Monitor reports

A list of reports available from data collected by the IMS Monitor, together with the principal performance data they contain, is shown in the following table.

The reports marked “MSC” in the list are only produced when MSC is active.

The order of the reports listed in the table matches the sequence of the output from the IMS Monitor Report Print Program. The duration of a monitoring snapshot might not include certain events necessary for an individual report, in which case only report headings or partial data are produced.

Table 75. Output sequence and information from IMS Monitor reports

| Report name                | Principal information                      |
|----------------------------|--------------------------------------------|
| System Configuration       | Monitor run documentation                  |
| Message Queue Pool         | Buffering and message I/O per transaction  |
| Message Format Buffer Pool | Count of I/Os                              |
| Latch Conflict Statistics  | IMS internal processing                    |
| General Wait Time Events   | Wait times for SNAPQ                       |
| Region and Jobname         | Monitor run documentation                  |
| Region Summary             | Elapsed times and count of DL/I calls (DC) |
| Region Wait                | Wait times                                 |
| Programs by Region         | Elapsed times for region usage             |
| Program Summary            | Overall program statistics                 |
| Communication Summary      | Elapsed times for lines                    |
| Communication Wait         | Wait times by line                         |
| Line Functions             | Count and size of blocks transmitted       |
| MSC Traffic (MSC)          | Count and routing of transactions          |

Table 75. Output sequence and information from IMS Monitor reports (continued)

| Report name               | Principal information                 |
|---------------------------|---------------------------------------|
| MSC Summaries (MSC)       | Count of transactions by destination  |
| MSC Queuing Summary (MSC) | Count and queuing time by link        |
| Transaction Queuing       | Queue loading statistics              |
| Reports                   | Count of space failures and deadlocks |
| Run Profile               | Monitor run documentation             |
| Distribution Appendix     | Event frequency distributions         |

The majority of the data items in IMS Monitor reports are elapsed times. These are normally expressed in microseconds. An entry of 1876534 is 1.876534 seconds or 1876 milliseconds. Any times that do not follow this convention show the unit of measure on the report.

You can also find counts of events under the heading OCCURRENCES, and figures that represent the number of bytes.

**Related concepts:**

“IMS Monitor reports for MSC” on page 763

## Documenting the monitoring run

For each trace interval there are several general reports or overall summaries of the processing that took place. You can use these reports as part of your IMS Monitor run documentation.

It is important to record as accurately as possible the conditions under which the trace was taken. Your documentation can include system status information obtained by the /DISPLAY command several times before and after the trace, an expected profile of the application program activity, and any desired processing events. The trace interval should represent typical processing loads and not be a biased or inadequate historical record.

### Adding to the System Configuration report data

The first general report titled SYSTEM CONFIGURATION is found under the page heading BUFFER POOL STATISTICS. It shows the modification level of the IMS and z/OS systems. You can add a list of IMS APARs applied and include the service levels of the application programs, especially if the latter are not permanent programs or are part of a staged implementation. The system configuration output is illustrated in the following section.

### Recording the Monitor trace interval

The heading of most IMS Monitor reports carries the trace start and stop times. It is shown in the format YEAR DAY (Julian) HH:MM:SS. The overall length of the trace interval is given in seconds under the heading TRACE ELAPSED TIME IN SECONDS. The following line shows how many trace records were placed on the IMS.MON data set. An example of the monitor trace interval recording is shown in the following example.

```
***IMS MONITOR***  BUFFER POOL STATISTICS  TRACE START 1993 130   5:55:15   TRACE STOP  1993 130   5:59:49   PAGE 0001
SYSTEM CONFIGURATION
```

```

SYSTEM CONFIGURATION      :
IMS VERSION                : 4
RELEASE LEVEL              :
MODIFICATION NUMBER       :

```

## Completing the Monitor run profile

A compact set of processing ratios is found at the end of the Run Profile report. The statistics summarize, for the monitor interval, the transaction throughput and the degree of DL/I and I/O activity. An example of the report is shown in the following example. In a DCCTL environment, DL/I activity is restricted to data communications calls and calls to GSAM databases. Database calls to other types of DL/I databases are not supported in DCCTL.

The lower part of the Run Profile report shows several ratios:

- Program elapsed time to DL/I elapsed time for each region
- DL/I elapsed time to wait time during DL/I processing
- Program elapsed time to other subsystem call elapsed time
- DL/I elapsed time to other subsystem call elapsed time

Each dependent region is identified by a sequence number, starting at region 1.

```

IMS MONITOR  **RUN PROFILE**      TRACE START 1993 130   5:55:15      TRACE STOP 1993 130   5:59:49  PAGE 0184
TRACE ELAPSED TIME IN SECONDS.....274.6
TOTAL NUMBER OF MESSAGES DEQUEUED....1403
TOTAL NUMBER OF SCHEDULES          .....173
NUMBER OF TRANSACTIONS PER SECOND   5.1
TOTAL NUMBER OF DL/I CALLS ISSUED...18632
NUMBER OF DL/I CALLS PER TRANSACTION 13.2
NUMBER OF OSAM BUFFER POOL I/O'S.    0,      0.0 PER TRANSACTION
NUMBER OF MESSAGE QUEUE POOL I/O'S.....0,    0.0 PER TRANSACTION
NUMBER OF FORMAT BUFFER POOL I/O'S.....0,    0.0 PER TRANSACTION
RATIO OF PROGRAM ELAPSED TO DL/I ELAPSED:
    REGION 1: 1.09
    REGION 2: 1.09
    REGION 3: 1.00
    REGION 4: 1.02
    REGION 5: 1.01
    REGION 6: 1.00
    REGION 7: 1.00
    REGION 8: 1.00
    REGION 9: 1.17
    REGION 10: 1.00
    REGION 11: 1.00
    REGION 49: 1.03
    REGION 50: 1.19
RATIO OF DL/I ELAPSED TO DL/I IWAIT:
    REGION 1: 325.65
    REGION 2: 73.49
    REGION 4: 100.35
    REGION 5: 85.76
    REGION 6: 82.99
    REGION 47: 95.64
    REGION 48: 45.93
    REGION 49: 9.22

```

You can match the regions to the z/OS job name using the Region and Jobname report. The job names correspond to the step names on the EXEC statements of all the dependent regions started by the operator before the trace was started. The region job names are included on the monitor output page with the heading GENERAL REPORTS, as illustrated in “Detecting checkpoint effects” on page 801.

There are some generalized processing ratios that are given at the end of several buffer pool statistics reports. You can include them in the documented profile of

the trace interval. These are not specific to one application or system resource but can be used as indicators of variation across a series of monitor runs. In DCCTL, the ratios are:

- All waits divided by the total number of transactions  
This value can be found on the Message Queue Pool Report in “Monitoring message queue handling” on page 801. This ratio indicates on a per transaction basis the physical I/O activity required to handle the message queuing function.
- The total prefetch I/Os + immediate fetch I/Os + directory I/Os divided by the total number of transactions  
This value also appears on the example shown in “Monitoring MFS activity” on page 800. This ratio indicates on a per transaction basis the physical I/O activity required to handle the MFS function during the trace period.

## Verifying IMS Monitor report occurrences

When you examine the output from the IMS Monitor Report Print program, do not assume that the presence of a report heading implies that appropriate data is listed. System definition options and utility control statements affect the content of the output as follows:

- The output does not include a Call Summary report unless a control statement specifies DLI.
- The output does not include a set of Distribution reports unless a control statement specifies DIS or DISTRIBUTION. The column headed DISTRIBUTION NUMBER that occurs on many of the reports contains cross-references to items included in the Distribution reports.
- The output consists of just a Call Summary report if a control statement specifies ONLY DLI.

Because many of the summary reports require system status to calculate the difference between start and end values, and this status is obtained during /TRACE SET OFF processing, the IMS Monitor execution must end before termination of the IMS control region. If the trace was not stopped properly, the following message is issued:

```
NO QUEUE BUFFER POOL TRACES AT END TIME ON MONITOR LOG TAPE
****QUEUE BUFFER POOL REPORT CANCELLED****
```

Similarly, other summary reports are not produced.

The section MESSAGE FORMAT BUFFER POOL is included only if your system definition specifies devices using Message Format Service (MFS).

If the source data used to formulate a particular IMS Monitor report, or a section of that report, has not been recorded by the IMS Monitor during the trace interval, the report contains only the headings.

---

## Monitoring activity in dependent regions

The IMS Monitor gathers timing information for every dependent region identified in the /trace command active during the trace interval. It records the total of the elapsed times for each event, the time for the longest event encountered, and the average time for all recorded events.

There are three major reports that display timings. The reports and a list of their content are:

### Region Summary Report

- Scheduling and termination
- Schedule end to first call
- Elapsed execution with separate summaries shown for:
  - DL/I calls
  - External subsystem service and command calls
  - External subsystem database access calls
  - Checkpoint processing
  - Region occupancy

### Region Wait

- Waits during scheduling and termination
- Waits during DL/I calls
- Waits during external subsystem calls
- Waits during checkpoint

### Programs by Region

- Elapsed execution
- Schedule end to first call

These three reports are illustrated in the following examples.

Activities in dependent regions are placed in five timing categories:

- Elapsed time for scheduling and termination  
The scheduling process includes many preparatory events such as block loading from an active IMS.ACBLIBA/B data set and obtaining ownership of the PSB. The time required to terminate the region activity after the application program ends is also included.
- Elapsed time from end of schedule to first call  
This time is reserved for application program initialization and housekeeping prior to an initial call (to the message queue, or an external subsystem) that marks the beginning of control program services. It is a measure of processing that is not repeated when multiple transactions are processed in a single scheduling.
- Program elapsed time, including all calls  
This time encompasses the major application program processing and is measured from the first call to the return to or exit from the program.
- Elapsed time performing DL/I calls  
This time includes all DL/I calls. Each DL/I call event is measured from the time of the call to the return to the application program.
- Elapsed time performing external subsystem calls  
This time includes all external subsystem calls. Each external subsystem event is measured from the time of the call to the return to IMS.

The following example shows a sample of the region summary report.

|                                   |   |                        |                   |                   |                               |                   |                             |                   |           |
|-----------------------------------|---|------------------------|-------------------|-------------------|-------------------------------|-------------------|-----------------------------|-------------------|-----------|
| IMS MONITOR                       |   | ****REGION SUMMARY**** |                   |                   | TRACE START 1993 130 5:55:15  |                   | TRACE STOP 1993 130 5:59:49 |                   | PAGE 0011 |
|                                   |   | (A)                    |                   |                   | (B)                           |                   |                             |                   |           |
|                                   |   | .....ELAPSED TIME..... |                   |                   | NOT IWAIT TIME(ELAPSED-IWAIT) |                   |                             | DISTRIBUTION      |           |
| OCCURRENCES                       |   | TOTAL                  | MEAN              | MAXIMUM           | TOTAL                         | MEAN              | MAXIMUM                     | NUMBER            |           |
| <u>SCHEDULING AND TERMINATION</u> |   | <u>          </u>      | <u>          </u> | <u>          </u> | <u>          </u>             | <u>          </u> | <u>          </u>           | <u>          </u> |           |
| **REGION                          | 5 | 5                      | 4146              | 829               | 948                           | 4146              | 829                         | 948               | 287A,B    |

|                        |    |        |             |           |           |             |         |           |              |
|------------------------|----|--------|-------------|-----------|-----------|-------------|---------|-----------|--------------|
| **REGION               | 6  | 7      | 6028        | 861       | 1067      | 6028        | 861     | 1067      | 214A,B       |
| **REGION               | 8  | 8      | 6847        | 855       | 1098      | 6847        | 855     | 1098      | 129A,B       |
| **REGION               | 10 | 7      | 9664        | 1380      | 3668      | 9664        | 1380    | 3668      | 272A,B       |
| **REGION               | 47 | 6      | 5482        | 913       | 1021      | 5482        | 913     | 1021      | 145A,B       |
| **REGION               | 49 | 3      | 2612        | 870       | 917       | 2612        | 870     | 917       | 443A,B       |
| **TOTALS               |    | 123    | 126042      | 1024      |           | 126042      | 1024    |           |              |
| SCHEDULE TO FIRST CALL |    |        |             |           |           |             |         |           |              |
| **REGION               | 1  | 1      | 15479797    | 15479797  | 15479797  |             |         |           | 555          |
| **REGION               | 2  | 1      | 22376350    | 22376350  | 22376350  |             |         |           | 564          |
| **REGION               | 3  | 1      | 15169488    | 15169488  | 15169488  |             |         |           | 578          |
| **REGION               | 4  | 1      | 48146258    | 48146258  | 48146258  |             |         |           | 584          |
| **REGION               | 48 | 1      | 795351      | 795351    | 795351    |             |         |           | 592          |
| **REGION               | 49 | 4      | 2960425     | 740106    | 2951746   |             |         |           | 442          |
| **REGION               | 50 | 1      | 15713464    | 15713464  | 15713464  |             |         |           | 575          |
| **TOTALS               |    | 168    | 514286738   | 3061230   |           |             |         |           |              |
| ELAPSED EXECUTION      |    |        |             |           |           |             |         |           |              |
| **REGION               | 1  | 1      | 290146255   | 290146255 | 290146255 |             |         |           | 1            |
| **REGION               | 2  | 1      | 252290108   | 252290108 | 252290108 |             |         |           | 2            |
| **REGION               | 3  | 1      | 259496970   | 259496970 | 259496970 |             |         |           | 3            |
| **REGION               | 4  | 1      | 322812716   | 322812716 | 322812716 |             |         |           | 4            |
| **REGION               | 48 | 1      | 273871107   | 273871107 | 273871107 |             |         |           | 48           |
| **REGION               | 49 | 4      | 271703421   | 67925855  | 155176058 |             |         |           | 49           |
| **REGION               | 50 | 1      | 290379922   | 290379922 | 290379922 |             |         |           | 50           |
| **TOTALS               |    | 173    | 14238540145 | 82303700  |           |             |         |           |              |
| DL/I CALLS             |    |        |             |           |           |             |         |           |              |
|                        |    |        |             |           |           |             |         |           | IWT/CALL (C) |
| **REGION               | 1  | 60     | 264626241   | 4410437   | 88981490  | 263813671   | 4396894 | 88970053  | 247A,B,C     |
| **REGION               | 2  | 223    | 230505269   | 1033655   | 61048758  | 227368742   | 1019590 | 61011153  | 237A,B,C     |
| **REGION               | 3  | 29     | 257704383   | 8886358   | 690000514 | 257704383   | 8886358 | 690000514 | 98A,B,C      |
| **REGION               | 4  | 792    | 313735347   | 396130    | 52439653  | 310609035   | 392183  | 52439653  | 180A,B,C     |
| **REGION               | 49 | 592    | 262886317   | 444064    | 30202068  | 234394017   | 395935  | 30159782  | 177A,B,C     |
| **REGION               | 50 | 36     | 242591451   | 6738651   | 48651260  | 242591451   | 6738651 | 48651260  | 289A,B,C     |
| **TOTALS               |    | 18632  | 12386905286 | 664818    |           | 12024562411 | 645371  |           |              |
| IDLE FOR INTENT        |    |        |             |           |           |             |         |           |              |
| NONE                   |    |        |             |           |           |             |         |           |              |
| CHECKPOINT             |    |        |             |           |           |             |         |           |              |
| NONE                   |    |        |             |           |           |             |         |           |              |
| REGION OCCUPANCY       |    |        |             |           |           |             |         |           |              |
| **REGION               | 1  | 100.0% |             |           |           |             |         |           |              |
| **REGION               | 2  | 100.0% |             |           |           |             |         |           |              |
| **REGION               | 3  | 100.0% |             |           |           |             |         |           |              |
| **REGION               | 4  | 100.0% |             |           |           |             |         |           |              |
| **REGION               | 48 | 100.0% |             |           |           |             |         |           |              |
| **REGION               | 49 | 100.0% |             |           |           |             |         |           |              |
| **REGION               | 50 | 100.0% |             |           |           |             |         |           |              |

The following example shows a sample of the region wait report.

|                          |                      |                      |         |                     |          |                           |
|--------------------------|----------------------|----------------------|---------|---------------------|----------|---------------------------|
| IMS MONITOR              | ****REGION IWAIT**** | TRACE START 1993 130 | 5:55:15 | TRACE STOP 1993 130 | 5:59:49  | PAGE 0023                 |
| .....IWAIT TIME.....     |                      |                      |         |                     |          |                           |
| **REGION                 | 5 OCCURRENCES        | TOTAL                | MEAN    | MAXIMUM             | FUNCTION | DISTRIBUTION              |
| SCHEDULING + TERMINATION |                      |                      |         |                     |          |                           |
| SUB-TOTAL                |                      |                      |         |                     |          |                           |
| TOTAL                    |                      |                      |         |                     |          |                           |
| DL/I CALLS               |                      |                      |         |                     |          |                           |
|                          |                      | 11                   | 181816  | 16528               | 24375    | DD=IMMSTR2A DBH 117       |
|                          |                      | 8                    | 112831  | 14103               | 17846    | DD=IMMSTR1A DBH 118       |
|                          |                      | 5                    | 85460   | 17092               | 33717    | DD=IMMSTR3A DBH 119       |
|                          |                      | 1                    | 1633    | 1633                | 1633     | INT=BMPFPE05 BLR-64BIT 14 |
|                          |                      | 1                    | 4242    | 4242                | 4242     | PSB=DDLTRN17 BLR-64BIT 20 |
|                          |                      | 5                    | 58420   | 11684               | 14643    | DD=IMINDEXA VBH 120       |
|                          |                      | 12                   | 173866  | 14488               | 22152    | DD=PRODCNTA VBH 121       |
|                          |                      | 3                    | 100576  | 33525               | 68373    | DD=IMMSTR2B DBH 428       |
|                          |                      | 1                    | 17921   | 17921               | 17921    | DD=IMMSTR3B DBH 429       |
|                          |                      | 1                    | 17195   | 17195               | 17195    | DD=IMMSTR1B DBH 430       |
|                          |                      | 1                    | 13577   | 13577               | 13577    | DD=IMINDEXB VBH 431       |
|                          |                      | 3                    | 49928   | 16642               | 20396    | DD=PRODCNTB VBH 432       |
|                          |                      | 4                    | 10973   | 2743                | 2787     | DD=ITEMACTB DBH 453       |
|                          |                      | 2                    | 37680   | 18840               | 27664    | DD=IAINDEXB VBH 454       |
|                          |                      | 49                   | 1500067 | 30613               | 138284   | DD=INVENTRA DBH 472       |
|                          |                      | 23                   | 345595  | 15025               | 27613    | DD=VENDORDA VBH 473       |
|                          |                      | 1                    | 342952  | 342952              | 342952   | PI=VENDORDA...1 498       |
|                          |                      | 1                    | 14612   | 14612               | 14612    | PI=VNSINDXA...1 499       |
|                          |                      | 6                    | 69203   | 11533               | 19492    | DD=VNSINDXA VBH 500       |
| TOTAL                    |                      | 136                  | 3132672 | 23034               |          |                           |



The following example shows a sample of the programs-by-region report.

| IMS MONITOR   |   | ****PROGRAMS BY REGION**** |                  |                   | TRACE START 1993 130 5:55:15 |                     | TRACE STOP 1993 130 5:59:49 |          | PAGE 0069              |
|---------------|---|----------------------------|------------------|-------------------|------------------------------|---------------------|-----------------------------|----------|------------------------|
|               |   | (A)                        |                  |                   | (B)                          |                     |                             |          |                        |
|               |   | OCCURRENCES                | ELAPSED<br>TOTAL | EXECUTION<br>MEAN | TIME<br>MAXIMUM              | SCHEDULING<br>TOTAL | END TO FIRST CALL<br>MEAN   | MAXIMUM  | DISTRIBUTION<br>NUMBER |
| **REGION      | 1 |                            |                  |                   |                              |                     |                             |          |                        |
|               | 1 |                            |                  |                   |                              |                     |                             |          |                        |
| PROGSC6D      |   | 1                          | 290146255        | 290146255         | 290146255                    | 15479797            | 15479797                    | 15479797 | 885A,B                 |
| REGION TOTALS |   | 1                          | 290146255        | 290146255         |                              | 15479797            | 15479797                    |          |                        |
| **REGION      | 2 |                            |                  |                   |                              |                     |                             |          |                        |
|               | 2 |                            |                  |                   |                              |                     |                             |          |                        |
| PROGIT8C      |   | 1                          | 252290108        | 252290108         | 252290108                    | 22376350            | 22376350                    | 22376350 | 889A,B                 |
| REGION TOTALS |   | 1                          | 252290108        | 252290108         |                              | 22376350            | 22376350                    |          |                        |
| **REGION      | 3 |                            |                  |                   |                              |                     |                             |          |                        |
|               | 3 |                            |                  |                   |                              |                     |                             |          |                        |
| PROGTS1C      |   | 1                          | 259496970        | 259496970         | 259496970                    | 15169488            | 15169488                    | 15169488 | 893A,B                 |
| REGION TOTALS |   | 1                          | 259496970        | 259496970         |                              | 15169488            | 15169488                    |          |                        |
| **REGION      | 4 |                            |                  |                   |                              |                     |                             |          |                        |
|               | 4 |                            |                  |                   |                              |                     |                             |          |                        |
| PROGPS3D      |   | 1                          | 322812716        | 322812716         | 322812716                    | 48146258            | 48146258                    | 48146258 | 897A,B                 |
| REGION TOTALS |   | 1                          | 322812716        | 322812716         |                              | 48146258            | 48146258                    |          |                        |
| **REGION      | 5 |                            |                  |                   |                              |                     |                             |          |                        |
|               | 5 |                            |                  |                   |                              |                     |                             |          |                        |
| PROGPS3A      |   | 2                          | 62893103         | 31446551          | 40693590                     | 5435                | 2717                        | 2862     | 901A,B                 |
| PROGTS1B      |   | 1                          | 61794787         | 61794787          | 61794787                     | 2790                | 2790                        | 2790     | 1271A,B                |
| PROGPS3B      |   | 1                          | 18294458         | 18294458          | 18294458                     | 3104                | 3104                        | 3104     | 1350A,B                |
| PROGIT2B      |   | 1                          | 36095342         | 36095342          | 36095342                     | 2731                | 2731                        | 2731     | 1363A,B                |
| PROGSC2A      |   | 1                          | 93902771         | 93902771          | 93902771                     | 1667791             | 1667791                     | 1667791  | 1401A,B                |
| REGION TOTALS |   | 6                          | 272980461        | 45496743          |                              | 1681851             | 280308                      |          |                        |
| **REGION      | 6 |                            |                  |                   |                              |                     |                             |          |                        |
|               | 6 |                            |                  |                   |                              |                     |                             |          |                        |
| PROGIT1B      |   | 2                          | 39000315         | 19500157          | 23703429                     | 5286                | 2643                        | 2801     | 905A,B                 |
| PROGTS1B      |   | 1                          | 34293636         | 34293636          | 34293636                     | 3136                | 3136                        | 3136     | 1207A,B                |
| PROGPS3A      |   | 1                          | 51887767         | 51887767          | 51887767                     | 2534                | 2534                        | 2534     | 1278A,B                |
| PROGPS3B      |   | 2                          | 67375031         | 33687515          | 40291430                     | 17210570            | 8605285                     | 17213287 | 1328A,B                |
| PROGIT8A      |   | 1                          | 69132416         | 69132416          | 69132416                     | 3291                | 3291                        | 3291     | 1359A,B                |
| PROGSC4A      |   | 1                          | 30165017         | 30165017          | 30165017                     | 2571                | 2571                        | 2571     | 1433A,B                |
| REGION TOTALS |   | 8                          | 291854182        | 36481772          |                              | 17193752            | 2149219                     |          |                        |
| **REGION      | 7 |                            |                  |                   |                              |                     |                             |          |                        |
|               | 7 |                            |                  |                   |                              |                     |                             |          |                        |
| PROGSC2B      |   | 1                          | 269618583        | 269618583         | 269618583                    | 5047875             | 5047875                     | 5047875  | 909A,B                 |
| REGION TOTALS |   | 1                          | 269618583        | 269618583         |                              | 5047875             | 5047875                     |          |                        |
| **REGION      | 8 |                            |                  |                   |                              |                     |                             |          |                        |
|               | 8 |                            |                  |                   |                              |                     |                             |          |                        |
| PROGIT8A      |   | 1                          | 5181039          | 5181039           | 5181039                      | 2928                | 2928                        | 2928     | 913A,B                 |
| PROGPS3A      |   | 1                          | 27304257         | 27304257          | 27304257                     | 3350                | 3350                        | 3350     | 1132A,B                |
| PROGSC4B      |   | 1                          | 37286872         | 37286872          | 37286872                     | 3009                | 3009                        | 3009     | 1255A,B                |
| PROGIT2A      |   | 1                          | 36902995         | 36902995          | 36902995                     | 2850                | 2850                        | 2850     | 1298A,B                |
| PROGIT1B      |   | 1                          | 30407479         | 30407479          | 30407479                     | 2565                | 2565                        | 2565     | 1336A,B                |
| PROGIT1A      |   | 3                          | 109875360        | 36625120          | 45190114                     | 4279008             | 1426336                     | 4272096  | 1357A,B                |
| PROGIT8B      |   | 1                          | 23405220         | 23405220          | 23405220                     | 2679                | 2679                        | 2679     | 1395A,B                |
| REGION TOTALS |   | 9                          | 270363222        | 30040358          |                              | 4296389             | 477376                      |          |                        |

## Examining the effects of checkpoints

Checkpoint processing can be initiated by the control program at a specified frequency determined by the number of records placed on the system log. Other checkpoints can be caused by operator commands.

The checkpoint line of the Region Summary report at the end of the region 0 summary shows the following:

- The number of system checkpoint taken during the monitor interval
- The elapsed times
- The not-wait times

The wait time experienced during checkpoints is reported at the end of the first region summary on the Region Wait report. You can detect delays for each combination of DD name and module code. Typical entries here are for the

message queue data sets and the restart data set. If an wait for storage is the cause, the entry under the FUNCTION column is STG.=, followed by the identification of the pool.

## Measuring region occupancy

A measure of region activity is the percentage of region occupancy. This is broadly the ratio of the elapsed time a region is performing processing to the trace interval.

The region occupancy time does not include those times when no messages are available, when the block loading is delayed, or when the PSB cannot be used. The last section in the Region Summary report lists all active regions for which timed events were collected and shows the calculated percentage region occupancies.

## Monitoring application program elapsed time

The IMS Monitor can record measurements of elapsed times for each transaction and scheduling of an application program. It does this during the monitored interval while other programs are executing concurrently.

Elapsed times are calculated from the start of the first DL/I (or other) call to the end of that program. You can distinguish between time spent in application code and in DL/I processing. The event intervals are illustrated in the following figure.




Figure 93. Elapsed Time event intervals

Within the elapsed time for a DL/I call, the wait time to obtain segment data is recorded separately. Similarly, within the elapsed time for an external subsystem call, the processing time in the external subsystem is recorded separately as the wait time. The application processing (A/P) time includes many kinds of subsidiary service beyond the machine cycles expended by the program object code—such as subroutine loading, I/O to z/OS data sets, and any overlay processing. If the program is waiting to be dispatched or requires paging before it

can use the real storage, these delays are also accounted for in application program processing. Because a program can execute many transactions for each schedule, the elapsed time from schedule to first call is recorded separately. This time covers the initialization performed by the application program and also includes the time for loading the program.

The elapsed times are given in the Program Summary report. The following example is a sample of the program summary report. Programs are identified by their PSB name on individual lines in the report. Each line gives a summary of the activity for that PSB during the measured interval. The total number of schedules, DL/I calls, transactions completed (dequeued), and waits for DL/I call I/O calls, the and external subsystem processing are given. The report line gives calculated average times for:

- Elapsed time per schedule
- Processor time per schedule
- Schedule to first DL/I call per schedule
- Elapsed time per transaction

Frequencies for calls per transaction, I/O waits per DL/I call, waits per external subsystem call, and transactions dequeued per schedule are also given. A TOTALS line summarizes all activity for the PSBs active during the monitored interval. (The PSB DUMMY line reconciles any incomplete scheduling caused by a region stopping during scheduling or for a program that experiences a pseudoabend.)

| IMS MONITOR |             | ****PROGRAM SUMMARY**** |       |             |            | TRACE START 1993 130 5:55:15 |                   |                  |            | TRACE STOP 1993 130 5:59:49 |                            |            |                      | PAGE 0075 |  |
|-------------|-------------|-------------------------|-------|-------------|------------|------------------------------|-------------------|------------------|------------|-----------------------------|----------------------------|------------|----------------------|-----------|--|
|             |             |                         |       |             |            | (A).....(B).....             |                   |                  |            | (A).....(B).....            |                            |            |                      |           |  |
| PSBNAME     | NO. SCHEDS. | TRANS. DEQ.             | CALLS | CALLS /TRAN | I/O IWAITS | I/O IWAITS /CALL             | TRAN. DEQD. /SCH. | CPU TIME /SCHED. | DISTR. NO. | ELAPSED TIME /SCHED.        | SCHED. TO 1ST CALL /SCHED. | DISTR. NO. | ELAPSED TIME /TRANS. |           |  |
| PROGSC6D    | 1           | 13                      | 60    | 4.6         | 46         | 0.7                          | 13.0              | 10010            | 884A,B     | 290146255                   | 15479797                   | 886A,B     | 22318942             |           |  |
| PROGIT8C    | 3           | 17                      | 225   | 13.2        | 166        | 0.7                          | 5.6               | 90592            | 888A,B     | 256617508                   | 73283259                   | 890A,B     | 45285442             |           |  |
| PROGTS1C    | 2           | 25                      | 47    | 1.8         | 0          | 0.0                          | 12.5              | 10010            | 892A,B     | 239190808                   | 7586234                    | 894A,B     | 19135264             |           |  |
| PROGSP3D    | 1           | 23                      | 792   | 34.4        | 182        | 0.2                          | 23.0              | 10010            | 896A,B     | 322812716                   | 48146258                   | 898A,B     | 14035335             |           |  |
| PROGSP3A    | 13          | 36                      | 1246  | 34.6        | 267        | 0.2                          | 2.7               | 49782            | 900A,B     | 32801812                    | 2228611                    | 902A,B     | 11845098             |           |  |
| PROGIT1B    | 11          | 21                      | 99    | 4.7         | 0          | 0.0                          | 1.9               | 6341             | 904A,B     | 23212388                    | 2036217                    | 906A,B     | 12158870             |           |  |
| PROGSC2B    | 7           | 155                     | 3068  | 19.7        | 1845       | 0.6                          | 22.1              | 346112           | 908A,B     | 93655514                    | 789390                     | 910A,B     | 42296603             |           |  |
| PROGIT8A    | 12          | 28                      | 434   | 15.5        | 293        | 0.6                          | 2.3               | 34350            | 912A,B     | 30196795                    | 1745815                    | 914A,B     | 12941483             |           |  |
| PROGSP2C    | 1           | 10                      | 179   | 17.9        | 205        | 1.1                          | 10.0              | 10010            | 916A,B     | 221024429                   | 53642029                   | 918A,B     | 22102442             |           |  |
| PROGTS1B    | 8           | 20                      | 54    | 2.7         | 0          | 0.0                          | 2.5               | 5447             | 920A,B     | 39943245                    | 2895                       | 922A,B     | 15977298             |           |  |
| PROGSP3C    | 1           | 14                      | 468   | 33.4        | 117        | 0.2                          | 14.0              | 10010            | 924A,B     | 310644485                   | 35978027                   | 926A,B     | 22188891             |           |  |
| PROGIT1C    | 1           | 9                       | 32    | 3.5         | 0          | 0.0                          | 9.0               | 10010            | 930A,B     | 304892631                   | 30226173                   | 932A,B     | 33876959             |           |  |
| PROGSC2C    | 1           | 9                       | 160   | 17.7        | 101        | 0.6                          | 9.0               | 10010            | 934A,B     | 296909110                   | 22242652                   | 936A,B     | 32989901             |           |  |
| PROGIT2B    | 8           | 21                      | 393   | 18.7        | 63         | 0.1                          | 2.6               | 21703            | 938A,B     | 35126671                    | 1798496                    | 940A,B     | 13381589             |           |  |
| PROGIT2C    | 6           | 17                      | 211   | 12.4        | 39         | 0.1                          | 2.8               | 13312            | 942A,B     | 288883508                   | 50698467                   | 944A,B     | 101958885            |           |  |
| PROGTS1D    | 2           | 26                      | 50    | 1.9         | 0          | 0.0                          | 13.0              | 10010            | 950A,B     | 284944505                   | 10613350                   | 952A,B     | 21918808             |           |  |
| PROGSP3B    | 8           | 22                      | 770   | 35.0        | 169        | 0.2                          | 2.7               | 35737            | 954A,B     | 38016279                    | 2149158                    | 956A,B     | 13824101             |           |  |
| PROGIT1A    | 11          | 24                      | 106   | 4.4         | 0          | 0.0                          | 2.1               | 7925             | 958A,B     | 30883486                    | 1935855                    | 960A,B     | 14154931             |           |  |
| PROGSC4A    | 9           | 163                     | 1775  | 10.8        | 5101       | 2.8                          | 18.1              | 235921           | 963A,B     | 62172947                    | 3011199                    | 965A,B     | 3432862              |           |  |
| PROGSC6C    | 1           | 10                      | 44    | 4.4         | 38         | 0.8                          | 10.0              | 10010            | 967A,B     | 228098334                   | 46568124                   | 969A,B     | 22809833             |           |  |
| PROGSP2B    | 11          | 28                      | 557   | 19.8        | 604        | 1.0                          | 2.5               | 35069            | 971A,B     | 33309266                    | 1181831                    | 973A,B     | 13085783             |           |  |
| PROGIT8D    | 1           | 12                      | 175   | 14.5        | 133        | 0.7                          | 12.0              | 10010            | 975A,B     | 253392289                   | 21274169                   | 977A,B     | 21116024             |           |  |
| PROGSC4C    | 1           | 10                      | 98    | 9.8         | 349        | 3.5                          | 10.0              | 10010            | 979A,B     | 248736332                   | 25930126                   | 981A,B     | 24873633             |           |  |
| PROGSC6A    | 7           | 157                     | 789   | 5.0         | 457        | 0.5                          | 22.4              | 11703            | 983A,B     | 73936039                    | 115979                     | 985A,B     | 3296511              |           |  |
| PROGIT2A    | 7           | 22                      | 430   | 19.5        | 71         | 0.1                          | 3.1               | 28529            | 987A,B     | 37905001                    | 2982                       | 989A,B     | 12060682             |           |  |
| PROGSC2D    | 1           | 15                      | 280   | 18.6        | 180        | 0.6                          | 15.0              | 10010            | 991A,B     | 316194222                   | 41527764                   | 993A,B     | 21079614             |           |  |
| PROGSP2A    | 6           | 25                      | 490   | 19.6        | 548        | 1.1                          | 4.1               | 43177            | 995A,B     | 58277945                    | 2467506                    | 997A,B     | 13986707             |           |  |
| PROGSC2A    | 5           | 121                     | 2363  | 19.5        | 1420       | 0.6                          | 24.2              | 276187           | 1001A,B    | 88906184                    | 6022954                    | 1003A,B    | 3673809              |           |  |
| PROGIT2D    | 1           | 20                      | 361   | 18.0        | 62         | 0.1                          | 20.0              | 10010            | 1005A,B    | 386092737                   | 111426279                  | 1007A,B    | 19304636             |           |  |
| PROGSC4B    | 10          | 131                     | 1421  | 10.8        | 4115       | 2.8                          | 13.1              | 617016           | 1011A,B    | 53826667                    | 2632409                    | 1013A,B    | 4108905              |           |  |
| PROGSC4D    | 1           | 19                      | 197   | 10.3        | 668        | 3.3                          | 19.0              | 10010            | 1020A,B    | 227999124                   | 46667334                   | 1022A,B    | 11999953             |           |  |
| PROGSP2D    | 1           | 13                      | 240   | 18.4        | 291        | 1.2                          | 13.0              | 10010            | 1025A,B    | 327602445                   | 52935987                   | 1027A,B    | 25200188             |           |  |
| PROGSC6B    | 5           | 140                     | 694   | 4.9         | 395        | 0.5                          | 28.0              | 16884            | 1032A,B    | 78994223                    | 3290769                    | 1034A,B    | 2821221              |           |  |
| PROGIT1D    | 1           | 10                      | 36    | 3.6         | 0          | 0.0                          | 10.0              | 10010            | 1041A,B    | 290379922                   | 15713464                   | 1043A,B    | 29037992             |           |  |
| PROGIT8B    | 8           | 17                      | 288   | 16.9        | 190        | 0.6                          | 2.1               | 33436            | 1259A,B    | 35223857                    | 2902                       | 1261A,B    | 16575932             |           |  |
| **TOTALS    | 173         | 1403                    | 18632 | 13.2        | 18115      | 0.9                          | 8.1               | 90328            |            | 82303700                    | 2972755                    |            | 10148638             |           |  |

You can use the Call Summary report to examine the detail of the call processing for each program, itemized by type or call and summarized for the monitor interval. An extract from the multi-page output is given in the following call summary report example. The calls using an I/O PCB are given first and

sub-totaled. Then, the total calls, of each type, against each database PCB and each external subsystem are listed. The PSB TOTAL line marks the end of data for each program.

| IMS MONITOR |          |                  | ****CALL SUMMARY**** |            |      | TRACE START 1993 130 5:55:15 |        |                 | TRACE STOP 1993 130 5:59:49 |                    |                    | PAGE 0186         |                    |
|-------------|----------|------------------|----------------------|------------|------|------------------------------|--------|-----------------|-----------------------------|--------------------|--------------------|-------------------|--------------------|
|             |          | CALL             |                      | LEV        | STAT |                              |        | (C)             | (A)                         |                    | (B)                |                   |                    |
| PSB NAME    | PCB NAME | NAME             | FUNC                 | NO.SEGMENT | CODE | CALLS                        | IWAITS | IWAITS/<br>CALL | ..ELAPSED<br>MEAN           | TIME...<br>MAXIMUM | .NOT IWAIT<br>MEAN | TIME..<br>MAXIMUM | DISTRIB.<br>NUMBER |
| PROGSC6B    | I/O PCB  | ISRT             | ( )                  |            |      | 138                          | 0      | 0.00            | 372                         | 1240               | 372                | 1240              | 598A,B,C           |
|             |          | GU               | ( )                  |            |      | 134                          | 133    | 0.99            | 2600917                     | 20974615           | 2587532            | 20962866          | 602A,B,C           |
|             |          | (GU)             | ( )                  |            |      | 3                            | 0      | 0.00            | 15                          | 16                 | 15                 | 16                | 716A,B,C           |
|             |          | ASRT             | ( )                  |            |      | 3                            | 0      | 0.00            | 330                         | 333                | 330                | 333               | 869A,B,C           |
|             |          | GU               | ( )                  |            | QC   | 2                            | 1      | 0.50            | 17639806                    | 21219588           | 17634776           | 21209529          | 870A,B,C           |
|             |          | I/O PCB SUBTOTAL |                      |            |      |                              |        |                 |                             |                    |                    |                   |                    |
|             |          |                  |                      |            |      | 280                          | 134    | 0.47            | 1370910                     |                    | 1364469            |                   |                    |
|             |          | PSB TOTAL        |                      |            |      |                              |        |                 |                             |                    |                    |                   |                    |
|             |          |                  |                      |            |      | 280                          | 134    | 0.47            | 1370910                     |                    | 1364469            |                   |                    |
|             |          |                  |                      |            |      |                              |        |                 |                             |                    |                    |                   |                    |
| PROGSC2A    | I/O PCB  | ISRT             | ( )                  |            |      | 118                          | 0      | 0.00            | 381                         | 1496               | 381                | 1496              | 603A,B,C           |
|             |          | GU               | ( )                  |            |      | 114                          | 284    | 2.49            | 3304809                     | 21784513           | 3164423            | 21664181          | 632A,B,C           |
|             |          | (GU)             | ( )                  |            |      | 2                            | 0      | 0.00            | 17                          | 18                 | 17                 | 18                | 781A,B,C           |
|             |          | ASRT             | ( )                  |            |      | 3                            | 0      | 0.00            | 367                         | 444                | 367                | 444               | 871A,B,C           |
|             |          | GU               | ( )                  |            | QC   | 2                            | 5      | 2.50            | 19931897                    | 20045206           | 19799530           | 19925277          | 872A,B,C           |
|             |          | I/O PCB SUBTOTAL |                      |            |      |                              |        |                 |                             |                    |                    |                   |                    |
|             |          |                  |                      |            |      | 239                          | 289    | 1.20            | 1743339                     |                    | 1675270            |                   |                    |
|             |          | PSB TOTAL        |                      |            |      |                              |        |                 |                             |                    |                    |                   |                    |
|             |          |                  |                      |            |      | 239                          | 289    | 1.20            | 1743339                     |                    | 1675270            |                   |                    |
|             |          |                  |                      |            |      |                              |        |                 |                             |                    |                    |                   |                    |
| PROGSC2D    | I/O PCB  | ISRT             | ( )                  |            |      | 14                           | 0      | 0.00            | 377                         | 621                | 377                | 621               | 608A,B,C           |
|             |          | GU               | ( )                  |            |      | 14                           | 36     | 2.57            | 22360408                    | 52048566           | 22221852           | 51901313          | 634A,B,C           |
|             |          | I/O PCB SUBTOTAL |                      |            |      |                              |        |                 |                             |                    |                    |                   |                    |
|             |          |                  |                      |            |      | 28                           | 36     | 1.28            | 11180393                    |                    | 11111115           |                   |                    |

## Monitoring I/O for application program DL/I calls

The IMS Monitor report shows the total number of I/O occurrences and the total time the occurrences took for each application program executed during a monitored interval. The Program I/O report gives these two totals for all PSBs active during the monitored interval and includes the detailed breakdown of the I/O wait time as it was incurred by each PCB used by the program.

The detail of the report reveals much of the contention experienced during application program processing. Each type of conflict and the number of times it occurred are recorded for each I/O PCB. The report shows the total wait time, the highest wait experienced, and the average time. Subtotals are given for each PCB under a PSB, and for all PCBs under each PSB.

The DDN/FUNC column lists the data set by DD name. The MODULE column uses a code to indicate the source of the contention. The types of conflicts and codes are shown as follows:

- **Message handling**

**Code Conflict**

**MFS** MFS format library directory

**PMM** Message format buffer pool space or control block I/O

**QMG** Message queue management

- **Scheduling**

**Code Conflict**

**BLR** Load/read from ACBLIB

**MSC** MPP region initialization

**SMN** Virtual storage management

For external subsystem calls, the elapsed time to complete the processing is considered wait time. The DDN/FUNC column indicates the external subsystem call function, as follows:

- **External subsystems**

| Code | Subsystem call function               |
|------|---------------------------------------|
| AB0  | ABORT                                 |
| CT0  | Create thread                         |
| D50  | Terminate identify or thread, signoff |
| D80  | INIT                                  |
| I30  | Identify, command, echo, terminate    |
| I30  | Identify, terminate subsystem         |
| I50  | INIT                                  |
| I60  | Resolve-in-doubt                      |
| PR0  | Subsystem-not-operational             |
| P10  | Commit prepare (Phase 1)              |
| P20  | Commit continue (Phase 2)             |
| SO0  | Signon                                |
| SI0  | Identify                              |

The following example shows as sample of the program I/O report.

| IMS MONITOR |          | ****PROGRAM I/O****  |          | TRACE START 1993 022 14:00:18 |         | TRACE STOP 1993 022 14:02:20 |        | PAGE 0088 |  |
|-------------|----------|----------------------|----------|-------------------------------|---------|------------------------------|--------|-----------|--|
|             |          | .....IWAIT TIME..... |          |                               |         |                              |        |           |  |
| PSBNAME     | PCB NAME | IWAITS               | TOTAL    | MEAN                          | MAXIMUM | DDN/FUNC                     | MODULE |           |  |
| PROGHR1A    | I/O PCB  | 122                  | 2341116  | 19189                         | 70795   | HOTELDBA                     | DBH    |           |  |
|             |          | 34                   | 24177936 | 711115                        | 3950160 | **W F I                      |        |           |  |
|             |          | 40                   | 23652665 | 591316                        | 2668917 | **W F I                      |        |           |  |
|             |          | 5                    | 67613    | 13522                         | 21214   | SHMSG                        | QMG    |           |  |
|             |          | 4                    | 110363   | 27590                         | 60486   | QBLKS                        | QMG    |           |  |
| PCB TOTAL   |          |                      |          |                               |         |                              |        |           |  |
|             |          | 131                  | 2519092  | 19229                         |         |                              |        |           |  |
| PSB TOTAL   |          |                      |          |                               |         |                              |        |           |  |
|             |          | 305                  | 6725063  | 20049                         |         |                              |        |           |  |
| PROGDE1A    | TRMNALDA | 20                   | 624677   | 31233                         | 68252   | TRMNALDA                     | VBH    |           |  |
|             |          | 1                    | 275811   | 275811                        | 275811  | PI TRMNALDA....              |        |           |  |
| PCB TOTAL   |          |                      |          |                               |         |                              |        |           |  |
|             |          | 21                   | 900488   | 42880                         |         |                              |        |           |  |
|             | I/O PCB  | 16                   | 488812   | 30550                         | 79980   | TRMNALDA                     | VBH    |           |  |
|             |          | 1                    | 16118    | 16118                         | 16118   | SHMSG                        | QMG    |           |  |
| PCB TOTAL   |          |                      |          |                               |         |                              |        |           |  |
|             |          | 17                   | 504930   | 29701                         |         |                              |        |           |  |
|             | TABLEDBA | 16                   | 290471   | 18154                         | 33254   | TABLEDA                      | DBH    |           |  |
| PCB TOTAL   |          |                      |          |                               |         |                              |        |           |  |
|             |          | 16                   | 290471   | 18154                         |         |                              |        |           |  |
| PSB TOTAL   |          |                      |          |                               |         |                              |        |           |  |
|             |          | 54                   | 1695889  | 31405                         |         |                              |        |           |  |
| PROGHR2B    | HOTELDBB | 8                    | 698384   | 87298                         | 184475  | HOTELDBB                     | DBH    |           |  |
|             |          | 4                    | 5820650  | 1455162                       | 1455278 | PI HOSINDEXB....             |        |           |  |
|             |          | 4                    | 4481024  | 1120256                       | 1209075 | PI HOTELDBB....              |        |           |  |
|             |          | 2                    | 260817   | 130408                        | 232750  | HOSINDEXB                    | VBH    |           |  |
|             |          | 7                    | 106623   | 15231                         | 16410   | HOSINDEXB                    | VBH    |           |  |
|             |          | 1                    | 15366    | 15366                         | 15366   | HOTELDBD                     | DBH    |           |  |
| PCB TOTAL   |          |                      |          |                               |         |                              |        |           |  |
|             |          | 26                   | 11382864 | 437802                        |         |                              |        |           |  |

```

PSB TOTAL
-----
PROGHR2A  HOTELDBA      26  11382864  437802
                        17    655801    38576    366108  HOSINDXA  VBH
                        73   1836721    25160    82141   HOTELDBA  DBH
                        2    54663    27331    41975   HOTELDBD  DBH
                        1     9887    9887    9887   HOTELDBC  DBH
                        2    851042   845635    845635  HOSINDOA  VBH

PCB TOTAL
-----
I/O PCB      95    3408114    35874
              20    575847    28792    74227   HOTELDBA  DBH
              21    370390    17637    43153   HOSINDXA  VBH
IMS MONITOR ****PROGRAM I/O**** TRACE START 1993 022 14:00:18 TRACE STOP 1993 022 14:02:20 PAGE 0089
.....IWAIT TIME.....
PSBNAME  PCB NAME      IWAITS      TOTAL      MEAN      MAXIMUM      DDN/FUNC  MODULE
-----  -
PROGHR2A  I/O PCB      5      4654544    930908    2020043    **W F I
              8      32796604   4099575    9328891    **W F I

PCB TOTAL
-----
PSB TOTAL      41      946237    23078
PROGPS2A  LOGIMA      136    4354351    32017
              89    2046670    22996    73593   IMMSTR3A  VBH
              612   53886417   88049    185674   IMMSTR1A  VBH
              3     44906    14968    20788   IMINDEXA  VBH

PCB TOTAL
-----
              704    55977993    79514
              469    11742900    25038    170337   COMPOSDA  DBH
              329    8198418    24919    91422   CPINDEXA  VBH

PCB TOTAL
-----
I/O PCB      798    19941318    24989
              3     47511    15837    20806   SHMSG      QMG

PCB TOTAL
-----
PSB TOTAL      3     47511    15837
PROGSC6C  I/O PCB      1505   75966822    50476
              52    2698602    51896    473763   INVENTRC  VBH
              4     70921    17730    34241   SHMSG      QMG
              3     50699    16899    24724   QBLKS      QMG

PCB TOTAL
-----
              59    2820222    47800
              55    2666884    48488    210752   INVENTRC  VBH
              50    797587    15951    41706   ININDEXC  VBH
              1    119253    119253    119253  PI INVENTRC...1
              1     8634    8634    8634   INVENTRB  VBH
              2     83947    41973    53936   INVENTRA  VBH

PCB TOTAL
-----
PSB TOTAL      109    3676305    33727
PROGHR2D  I/O PCB      168    6496527    38669
              21    2285296    108823    199223   HOTELDBD  DBH
              28    762370    27227    111860   HOSINDXD  VBH
              1    11685    11685    11685   SHMSG      QMG

PCB TOTAL
-----
HOTELDBD      50    3059351    61187
              96    6279107    65407    139032   HOTELDBD  DBH
MONITOR ****PROGRAM I/O**** TRACE START 1993 022 14:00:18 TRACE STOP 1993 022 14:02:20 PAGE 0090
.....IWAIT TIME.....
PSBNAME  PCB NAME      IWAITS      TOTAL      MEAN      MAXIMUM      DDN/FUNC  MODULE
-----  -
PROGHR2D  HOTELDBD      31    2130585    68728    769130   HOSINDXD  VBH
              3    115999    38666    56394   HOTELDBA  DBH
              2    69833    34916    43470   HOTELDBC  DBH
              2    41430    20715    28020   HOSINDOD  VBH
              4    5515374   1378843   1458884  PI HOSINDXD....
              4    3997017    999254   1026228  PI HOTELDBD....

```

|     |       |     |                 |
|-----|-------|-----|-----------------|
| PCB | TOTAL |     |                 |
|     |       | 142 | 18149345 127812 |
| PSB | TOTAL |     |                 |
|     |       | 192 | 21208696 110461 |

The I/O waits for the calls to the I/O PCB are grouped as the first entries for a PSB. For DC DL/I calls, the data set for which the I/O took place is indicated under the DDN/FUNC heading, and the module code indicates what type of conflict caused the wait. For external subsystem calls, the function is indicated under the DDN/FUNC heading, and the module code indicates the source of the call entry.

Names other than LGMSG and SHMSG can appear under the DDN/FUNC column for I/O PCBs.

If the program is designated as wait for input and has to wait for the input of the next message, the wait entry is marked \*\*WFI under the DDN/FUNC heading and no entry appears in the MODULE column. The time spent waiting for the next input message is shown under wait time. \*\*WFI entries are shown for information only and their values are not used to compute statistics.

## Monitoring MFS activity

You can obtain a summary of all activity that occurs for management of message format buffer pool use from the Message Format Buffer Pool report.

The report is illustrated in the following example. The data shows the counts at the start and end of the trace interval and their difference.

When message formatting occurs, the appropriate message blocks must reside in the message format buffer pool, a DIF/MID pair for input or a DOF/MOD pair for output. If the blocks are not already in the buffer, I/O to the active IMSFORMAT/A/B library must occur. Block retrieval can involve a prior directory lookup, or be direct, using an index kept in the pool.

Many of the counts reveal details of internal event management. The number of times there is no directory entry for a block implies extra directory lookup I/O. Delays caused by unavailable FRE entries are recorded as request-ignored counts.

\*\*\*IMS MONITOR\*\*\* BUFFER POOL STATISTICS TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0007

### M E S S A G E F O R M A T B U F F E R P O O L

|                                                            | 5:55:15<br>START TRACE | 5:59:49<br>END TRACE | DIFFERENCE |
|------------------------------------------------------------|------------------------|----------------------|------------|
| NUMBER OF P/F REQUESTS                                     | 0                      | 0                    | 0          |
| NUMBER OF I/F REQUESTS                                     | 18                     | 20                   | 2          |
| NUMBER OF I/F I/O'S                                        | 2                      | 2                    | 0          |
| NUMBER OF TIMES POOL COMPRESS WOULD BE SUCCESSFUL          | 0                      | 0                    | 0          |
| NUMBER OF DIRECTORY I/O OPERATIONS                         | 2                      | 2                    | 0          |
| NUMBER OF TIMES BLOCK WASHED FOR FRE                       | 0                      | 0                    | 0          |
| NUMBER OF TIMES P/F REQUEST IGNORED                        | 0                      | 0                    | 0          |
| NUMBER OF F/B REQUESTS                                     | 18                     | 20                   | 2          |
| NUMBER OF TIMES F/B REQUEST IGNORED                        | 0                      | 0                    | 0          |
| NUMBER OF TIMES I/F ON F/B QUEUE                           | 16                     | 18                   | 2          |
| NUMBER OF TIMES I/F ON I/F QUEUE                           | 0                      | 0                    | 0          |
| NUMBER OF TIMES F/B ON I/F QUEUE                           | 18                     | 20                   | 2          |
| NUMBER OF TIMES P/F ON I/F QUEUE                           | 0                      | 0                    | 0          |
| NUMBER OF TIMES P/F ON F/B QUEUE                           | 0                      | 0                    | 0          |
| NUMBER OF TIMES THERE WAS NO DIR ENTR FOR A BLOCK          | 0                      | 0                    | 0          |
| NUMBER OF TIMES I/O ERRORS POINT OR READ MACRO             | 0                      | 0                    | 0          |
| NUMBER OF IMMEDIATE I/O REQUESTS WAITED DUE TO MAXIMUM I/O | 0                      | 0                    | 0          |
| NUMBER OF REQUESTS SATISFIED BY INDEX/DYNAMIC DIRECTORY    | 0                      | 0                    | 0          |

QUOTIENT : IMMEDIATE FETCH I/O'S + DIRECTORY I/O'S OPERATIONS = 0.00

TOTAL NUMBER OF TRANSACTIONS

## Monitoring message queue handling

A key resource that directly affects the efficiency of transaction processing is the message queue pool and the management of the I/O to the message queues. You can examine the activity by looking at the Message Queue Pool report.

The following example illustrates the message queue pool report contents. Counts of activities are given at start and end of the trace interval and as the differences between start and end numbers.

\*\*\*IMS MONITOR\*\*\* BUFFER POOL STATISTICS TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0002

### MESSAGE QUEUE POOL

|                                                     | 5:55:15<br>START TRACE | 5:59:49<br>END TRACE | DIFFERENCE |
|-----------------------------------------------------|------------------------|----------------------|------------|
| NUMBER OF LOCATE CALLS FROM QMGR                    | 54204                  | 68436                | 14232      |
| NUMBER OF RECORD RELEASE CALLS FROM QMGR            | 16431                  | 20738                | 4307       |
| NUMBER OF LOCATE AND ALTER CALLS FROM QMGR          | 131593                 | 164744               | 33151      |
| NUMBER OF REQUESTS TO PURGE THE Q POOL              | 2                      | 2                    | 0          |
| NUMBER OF ADDRESS TO DRRN TRANSLATION REQUESTS      | 21351                  | 27076                | 5725       |
| NUMBER OF REQUESTS TO WAIT FROM QMGR                | 0                      | 0                    | 0          |
| NUMBER OF READ REQUESTS                             | 962                    | 962                  | 0          |
| NUMBER OF WRITE REQUESTS(TOTAL)                     | 499                    | 499                  | 0          |
| NUMBER OF WRITES DONE BY PURGE                      | 499                    | 499                  | 0          |
| NUMBER OF WAITS FOR PURGE COMPLETION                | 1                      | 1                    | 0          |
| NUMBER OF WAITS BECAUSE NO BUFFER AVAILABLE         | 0                      | 0                    | 0          |
| NUMBER OF WAITS FOR OTHER DECB TO READ THIS BUFFER  | 823                    | 823                  | 0          |
| NUMBER OF WAITS FOR OTHER DECB TO WRITE THIS BUFFER | 0                      | 0                    | 0          |
| NUMBER OF WAITS FOR CONFLICTING END DEQ BUFFER REQ  | 0                      | 0                    | 0          |
| NUMBER OF PSBS UNCHAINED FROM BUFFERS               | 0                      | 0                    | 0          |
| NUMBER OF CALLS TO QMGR.(TOTAL)                     | 48164                  | 62213                | 14049      |
| NUMBER OF CALLS TO REPOSITION A LOST BUFFER         | 0                      | 0                    | 0          |
| NUMBER OF CALLS TO ENQ A MESSAGE                    | 10583                  | 13441                | 2858       |
| NUMBER OF CALLS TO DEQ ONE OR MORE MESSAGE          | 6321                   | 7767                 | 1446       |
| NUMBER OF CALLS TO CANCEL INPUT OR OUTPUT           | 119                    | 121                  | 2          |

QUOTIENT : TOTAL NUMBER OF OSAM READS + OSAM WRITES + ALL IWAITS = 0.00

TOTAL NUMBER OF TRANSACTIONS

## Detecting checkpoint effects

When a checkpoint command specifies SNAPQ, the current status of all message queues is written to the system log. This prevents any message handling on behalf of queue management.

General Iwait Time Events report records the wait time incurred by the SNAPQ. The following example shows the activity on the summary line QMGR SNAPQ CHECK. The number of occurrences is given with the total, average, and maximum wait times.

IMS MONITOR \*\* GENERAL REPORTS \*\* TRACE START 1993 130 5:55:15 TRACE STOP 1993 130 5:59:49 PAGE 0009  
GENERAL IWAIT TIME EVENTS

| EVENT            | OCURRENCES | TOTAL | MEAN | MAXIMUM | DISTRIBUTION<br>NUMBER |
|------------------|------------|-------|------|---------|------------------------|
| QMGR SNAPQ CHECK | 0          | 0     | 0    | 0       | 0                      |

REGION AND JOBNAME REPORT

| REG. NO. | JOB NAME |
|----------|----------|
| 1        | MPR1A100 |
| 2        | MPR1A209 |
| 3        | MPR1A210 |
| 4        | MPR1A211 |
| 5        | MPR1A103 |
| 6        | MPR1A101 |



7 MPR1A115  
 8 MPR1A116  
 9 MPR1A216  
 10 MPR1A200  
 11 MPR1A217  
 12 MPR1A119  
 13 MPR1A218  
 14 MPR1A219  
 15 MPR1A104  
 16 MPR1A220  
 17 MPR1A203  
 18 MPR1A123  
 19 MPR1A222  
 20 MPR1A105  
 21 MPR1A124  
 22 MPR1A223  
 23 MPR1A107  
 24 MPR1A224  
 25 MPR1A106  
 26 MPR1A206  
 27 MPR1A205  
 28 MPR1A108  
 29 MPR1A109  
 30 MPR1A208  
 31 MPR1A111  
 32 MPR1A112  
 33 MPR1A113  
 34 MPR1A204  
 35 MPR1A114  
 36 MPR1A102  
 48 MPR1A121  
 49 MPR1A122  
 50 MPR1A221

## Transaction Queuing report

In addition to monitoring the efficiency of message handling, you can monitor the service provided for each application, by looking at the size of the transaction queues at each scheduling of their processing programs.

The Transaction Queuing report shown in the following example records, for each transaction, the minimum, average, and maximum counts at scheduling time. The total number of dequeued transactions (or transactions that have been fully processed) during the monitored interval is given for each transaction code. The average number of transactions processed for each scheduling is given in the DEQUEUED MEAN column.

| IMS MONITOR | ****TRANSACTION QUEUING**** |                |                       | TRACE START                   | 1993 130         | 5:55:15               | TRACE STOP             | 1993 130 | 5:59:49 | PAGE 0181 |
|-------------|-----------------------------|----------------|-----------------------|-------------------------------|------------------|-----------------------|------------------------|----------|---------|-----------|
| TRANSACTION | NUMBER DEQUED               | NUMBER SCHEDS. | ..ON QUEUE<br>MINIMUM | (B)<br>WHEN SCHEDULED<br>MEAN | .....<br>MAXIMUM | (A)<br>DEQUED<br>MEAN | DISTRIBUTION<br>NUMBER |          |         |           |
| SC6X        | 13                          | 1              | 0                     | 0.00                          | 0                | 13.00                 | 883A,B                 |          |         |           |
| IT8W        | 17                          | 3              | 0                     | 0.00                          | 0                | 5.66                  | 887A,B                 |          |         |           |
| TS1Z        | 16                          | 1              | 0                     | 0.00                          | 0                | 16.00                 | 891A,B                 |          |         |           |
| PS3X        | 23                          | 1              | 0                     | 0.00                          | 0                | 23.00                 | 895A,B                 |          |         |           |
| PS3Y        | 17                          | 7              | 0                     | 0.00                          | 0                | 2.42                  | 899A,B                 |          |         |           |
| IT1V        | 11                          | 6              | 0                     | 0.00                          | 0                | 1.83                  | 903A,B                 |          |         |           |
| SC2Z        | 143                         | 2              | 0                     | 0.00                          | 0                | 71.50                 | 907A,B                 |          |         |           |
| IT8U        | 12                          | 7              | 0                     | 0.00                          | 0                | 1.71                  | 911A,B                 |          |         |           |
| PS2W        | 10                          | 1              | 0                     | 0.00                          | 0                | 10.00                 | 915A,B                 |          |         |           |
| TS1U        | 12                          | 4              | 0                     | 0.00                          | 0                | 3.00                  | 919A,B                 |          |         |           |
| PS3W        | 14                          | 1              | 0                     | 0.00                          | 0                | 14.00                 | 923A,B                 |          |         |           |
| IT8Y        | 16                          | 5              | 0                     | 0.00                          | 0                | 3.20                  | 927A,B                 |          |         |           |
| IT1W        | 9                           | 1              | 0                     | 0.00                          | 0                | 9.00                  | 929A,B                 |          |         |           |
| SC2W        | 9                           | 1              | 0                     | 0.00                          | 0                | 9.00                  | 933A,B                 |          |         |           |
| IT2V        | 13                          | 5              | 0                     | 0.00                          | 0                | 2.60                  | 937A,B                 |          |         |           |
| IT2W        | 17                          | 6              | 0                     | 0.00                          | 0                | 2.83                  | 941A,B                 |          |         |           |
| TS1V        | 9                           | 1              | 0                     | 0.00                          | 0                | 9.00                  | 945A,B                 |          |         |           |
| SC2V        | 12                          | 5              | 0                     | 0.00                          | 0                | 2.40                  | 947A,B                 |          |         |           |
| TS1W        | 11                          | 1              | 0                     | 0.00                          | 0                | 11.00                 | 949A,B                 |          |         |           |
| PS3V        | 13                          | 3              | 0                     | 0.00                          | 0                | 4.33                  | 953A,B                 |          |         |           |
| IT1U        | 9                           | 6              | 0                     | 0.00                          | 0                | 1.50                  | 957A,B                 |          |         |           |
| SC4U        | 11                          | 5              | 0                     | 0.00                          | 0                | 2.20                  | 962A,B                 |          |         |           |

|      |     |   |   |      |   |        |         |
|------|-----|---|---|------|---|--------|---------|
| SC6W | 10  | 1 | 0 | 0.00 | 0 | 10.00  | 966A,B  |
| PS2V | 8   | 6 | 0 | 0.00 | 0 | 1.33   | 970A,B  |
| IT8X | 12  | 1 | 0 | 0.00 | 0 | 12.00  | 974A,B  |
| SC4W | 10  | 1 | 0 | 0.00 | 0 | 10.00  | 978A,B  |
| SC6U | 14  | 6 | 0 | 0.00 | 0 | 2.33   | 982A,B  |
| IT2Y | 9   | 3 | 0 | 0.00 | 0 | 3.00   | 986A,B  |
| SC2X | 15  | 1 | 0 | 0.00 | 0 | 15.00  | 990A,B  |
| PS2Y | 17  | 2 | 0 | 0.00 | 0 | 8.50   | 994A,B  |
| SC4Y | 152 | 4 | 0 | 0.50 | 1 | 38.00  | 998A,B  |
| SC2Y | 106 | 2 | 0 | 0.00 | 0 | 53.00  | 1000A,B |
| IT2X | 20  | 1 | 0 | 0.00 | 0 | 20.00  | 1004A,B |
| SC2U | 15  | 3 | 0 | 0.00 | 0 | 5.00   | 1008A,B |
| SC4Z | 123 | 5 | 0 | 0.60 | 1 | 24.60  | 1010A,B |
| TS1X | 15  | 1 | 0 | 0.00 | 0 | 15.00  | 1015A,B |
| SC4X | 19  | 1 | 0 | 0.00 | 0 | 19.00  | 1019A,B |
| PS2X | 13  | 1 | 0 | 0.00 | 0 | 13.00  | 1024A,B |
| PS2Z | 20  | 5 | 0 | 0.00 | 0 | 4.00   | 1028A,B |
| SC6Z | 130 | 1 | 0 | 0.00 | 0 | 130.00 | 1031A,B |
| SC6V | 10  | 4 | 0 | 0.00 | 0 | 2.50   | 1035A,B |
| SC6Y | 143 | 1 | 0 | 0.00 | 0 | 143.00 | 1037A,B |
| IT1X | 10  | 1 | 0 | 0.00 | 0 | 10.00  | 1040A,B |
| PS3U | 19  | 6 | 0 | 0.00 | 0 | 3.16   | 1131A,B |
| IT2U | 13  | 4 | 0 | 0.00 | 0 | 3.25   | 1146A,B |

## Monitoring line activity

You can obtain a summary of all occurrences of activity for each node that handles message traffic during the monitored interval. The elapsed times and NOT-WAIT times are given in categories of total, mean, and maximum times for each communication line in the Communication Summary report.

The following example shows a sample of the communication summary report.

**Requirement:** You must match which physical devices are using the line to the Stage 1 output from system definition. The line numbers are assigned sequentially, according to their physical occurrence in the Stage 1 input deck.

If your online system specifies the prefetch option for MFS blocks in the control region JCL, the last line of the report contains the statistics for all prefetch events.

|                        |                               |                                 |          |         |                                        |          |         |                        |
|------------------------|-------------------------------|---------------------------------|----------|---------|----------------------------------------|----------|---------|------------------------|
| IMS MONITOR            | ****COMMUNICATION SUMMARY**** | TRACE START                     | 1993 130 | 5:55:15 | TRACE STOP                             | 1993 130 | 5:59:49 | PAGE 0089              |
|                        |                               | (A)                             |          |         | (B)                                    |          |         |                        |
| NODE OR<br>LINE NUMBER | OCCURRENCES                   | .....ELAPSED TIME.....<br>TOTAL | MEAN     | MAXIMUM | NOT IWAIT TIME(ELAPSED-IWAIT)<br>TOTAL | MEAN     | MAXIMUM | DISTRIBUTION<br>NUMBER |
| PMT01A                 | 3                             | 2396                            | 798      | 1547    | 2396                                   | 798      | 1547    | 1467A,B                |
| 19                     | 182                           | 92155                           | 506      | 1106    | 92155                                  | 506      | 1106    | 1493A,B                |
| 2                      | 59                            | 2280                            | 38       | 41      | 2280                                   | 38       | 41      | 1515A,B                |
| TOTAL                  |                               |                                 |          |         |                                        |          |         |                        |
|                        | 244                           | 96831                           | 396      |         | 96831                                  | 396      |         |                        |

You can also investigate the amount of data transmitted across nodes with the Line Functions report. The following example shows a sample of a line-function report. The report distinguishes between input data and output data. The number of blocks of data and the average and maximum size of the blocks are recorded for data received by IMS and for transmitted data.

This report also includes a measure of how inactive the lines are. An inactive interval is assumed to be the difference between the time that marks the end of the last input block received and the starting time for output transmission. These occurrences of inactivity are termed turnaround intervals, and the report cumulates the number of occurrences as well as the average and maximum times associated with these intervals.

If the line is being used by an MFS supported terminal, a count of the number of requests for next page for a multi-page message is recorded.

If link traffic for coupled multiple systems is recorded, a set of three reports follows the Line Functions report.

| IMS MONITOR            |                | ****LINE FUNCTIONS*** |         |         | TRACE START 1993 130 5:55:15 |         |         |         | TRACE STOP 1993 130 5:59:49 |          |          |          | PAGE 0091          |       |
|------------------------|----------------|-----------------------|---------|---------|------------------------------|---------|---------|---------|-----------------------------|----------|----------|----------|--------------------|-------|
|                        |                | (A).....(B).....      |         |         |                              |         |         |         |                             |          |          |          |                    |       |
| NODE OR<br>LINE NUMBER | DEVICE<br>TYPE | RECEIVE               | RECEIVE | RECEIVE | TRANS.                       | TRANS.  | TRANS.  | DIST.   | TURN<br>AROUND<br>INTERVALS | MEAN     | MAX.     | DIST.    | PAGING<br>REQUESTS |       |
|                        |                | BLOCKS                | BLKSIZE | BLKSIZE | BLOCKS                       | BLKSIZE | BLKSIZE | NUMBER  |                             | INTERVAL | INTERVAL | INTERVAL |                    | NUMB. |
| PMT01A                 | 3270V          | 1                     | 29      | 29      | 2                            | 170     | 171     | 1468A,B | 3                           | 798      |          | 1547     | 1466               | 0     |
| 19                     | XXXX           |                       |         |         |                              |         |         |         | 182                         | 506      |          | 1106     | 1492               | 0     |
| 2                      | LOC SYS        |                       |         |         |                              |         |         |         | 59                          | 38       |          | 41       | 1514               | 0     |
| ALL LINES              |                |                       |         |         |                              |         |         |         |                             |          |          |          |                    |       |
|                        |                | 1                     | 29      |         | 2                            | 170     |         |         | 244                         | 396      |          |          |                    | 0     |

Related concepts:

“IMS Monitor reports for MSC” on page 763

## Monitoring message handling efficiency

The IMS Monitor produces both summary and detailed information on asynchronous processing in the IMS control region. The arrival of data transmitted from terminals triggers the processing.

Application program responses also result in processing. The space in four major buffer pools and access to format, SPA, and message queue data sets are managed for the total communications traffic. Wait times are recorded when contention for pool space or I/O interrupts the processing of any of the communication tasks triggered by line activity. This information is contained in the Communication Wait report. The following example illustrates this report.

This report is complementary to the Communication Summary report in that the line number is used as an identification for the series of communication processing tasks.

|              |             |                             |      |         |                              |         |                             |     |           |  |
|--------------|-------------|-----------------------------|------|---------|------------------------------|---------|-----------------------------|-----|-----------|--|
| IMS MONITOR  |             | ****COMMUNICATION IWAIT**** |      |         | TRACE START 1993 130 5:55:15 |         | TRACE STOP 1993 130 5:59:49 |     | PAGE 0090 |  |
| NODE OR      |             | .....IWAIT TIME.....        |      |         |                              |         |                             |     | DIST.     |  |
| LINE NUMBER  | OCCURRENCES | TOTAL                       | MEAN | MAXIMUM | FUNCTION                     | BLKSIZE | MODULE                      | NO. |           |  |
| ALL LINES... |             |                             |      |         |                              |         |                             |     |           |  |
| PREFETCH I/O |             |                             |      |         |                              |         |                             |     |           |  |
|              | NONE        |                             |      |         |                              |         |                             |     |           |  |

## IMS internal resource usage

There are several summary reports that you can use to examine the level of internal contention for resources.

The following list gives a brief explanation of these reports.

### Pool Space Failure Summary report

The Pool Space Failure Summary report gives the number of times in each region a given amount of storage was unavailable. It shows the number of bytes and the identification of the pool as well as the number of occurrences of this failure to obtain storage. You can use this summary to determine whether you need to increase the buffer pool allocation by a system definition change or by overriding the number of buffers in the EXEC statements in the JCL.

The format of the report is shown in the following example.

#### POOL SPACE FAILURE SUMMARY

|       | POOL ID | BYTES REQ. | OCCURRENCES |
|-------|---------|------------|-------------|
|       | DLMP    | 8888       | 1           |
| TOTAL |         |            | 1           |

### Latch Conflict Statistics report

The basic serialization of the task processing in IMS is controlled by ownership of an IMS latch. When different programs are executing, they compete for the ownership. If they wait for the resource, the one possessing the latch has to post the other ITASK waiting for it. You can judge the level of contention for a resource and then investigate a set of changes to relieve the pressure.

The different types of latches and the counters that exhibit the level of contention are given in the Latch Conflict Statistics report. The following figure is an example of this report. The entries are organized according to the latch names.

When a system checkpoint is taken during the time the monitor is active, latch conflict statistics are reset to zero, thus corrupting the values presented in this report. If this situation exists, the following message will be inserted at the top of the report:

```
**** A CHECKPOINT OCCURRED DURING MONITOR RUN ****
**** LATCH CONFLICT STATISTICS ARE INVALID ****
**** SEE UTILITIES REFERENCE MANUAL ****
```

However, if the master terminal operator issues the /CHECKPOINT command with the STATISTICS keyword parameter, latch conflict statistics are reset to zero, but the IMS monitor is not notified. Therefore, DFSUTR20 cannot detect that the statistics have been corrupted and does not issue this message.

**Recommendation:** Do not issue statistics checkpoints while the Monitor is running.

The counters are primarily concerned with storage management and logging services. The statistics recorded are the number of times contentions occur, that is, the resource waits for a latch.

```
IMS MONITOR ** GENERAL REPORTS ** TRACE START 1993 209...
```

#### LATCH CONFLICT STATISTICS

| LATCH<br>NAMES | COUNT<br>FIELD | AT<br>START | AT<br>END | DIFF. |
|----------------|----------------|-------------|-----------|-------|
| LOGL           | CONTENTIONS    | 0           | 0         | 0     |
| SMGT           | CONTENTIONS    | 0           | 0         | 0     |
| XCNQ           | CONTENTIONS    | 0           | 0         | 0     |
| ACTL           | CONTENTIONS    | 0           | 0         | 0     |
| CBTS           | CONTENTIONS    | 0           | 0         | 0     |
| DBLK           | CONTENTIONS    | 0           | 0         | 0     |

---

## Using frequency distributions from IMS Monitor output

The reports that are derived from the IMS Monitor data records contain many summary lines where the mean time is given. If you are interested in the distribution of those timed events, rather than just average and maximum times, you can request the Report Print utility to individually record the events in a frequency distribution across a range of intervals.

Some distributions are not time dependent, such as those for transaction queue loads or transmitted block sizes.

The following tables show the major IMS Monitor reports and the type of frequency distributions generated for each report. Each type results in several distributions, depending on how many entries are in each section of the report. For each type of frequency distribution the data is cumulated in suitable intervals or ranges. The set of ranges used for each type is given an identifier, shown in the ID column.

- The following table shows the report distributions sorted by Region Summary.

*Table 76. Report distributions by region summary*

| Report name                       | ID  | Description   |
|-----------------------------------|-----|---------------|
| Scheduling and Termination        | D1  | Elapsed time  |
|                                   | D2  | Not wait time |
| Schedule end to 1st DL/I call     | D3  | N/A           |
| Elapsed execution time DL/I calls | D4  | N/A           |
|                                   | D5  | Elapsed time  |
| External Subsystem calls          | D6  | Not wait time |
| Waits per DL/I call               | D43 | Elapsed time  |
| Idle for intent Checkpoint        | D7  | N/A           |
|                                   | D8  | N/A           |
|                                   | D20 | Elapsed time  |
|                                   | D21 | Not wait time |

- The following table shows the report distributions Programs Region.

*Table 77. Report distributions by program region*

| Report name                   | ID  | Description |
|-------------------------------|-----|-------------|
| Elapsed execution time        | D30 | N/A         |
| Schedule and to 1st DL/I call | D31 | N/A         |

- The following table shows the report distributions sorted by Program Summary.

*Table 78. Report distributions by program summary*

| Report name                        | ID  | Description |
|------------------------------------|-----|-------------|
| Processor time per schedule        | D15 | N/A         |
| Transactions dequeued per schedule | D14 | N/A         |
| Elapsed time per schedule          | D9  | N/A         |
| Schedule end to 1st DL/I call      | D10 | N/A         |

- The following table shows the report distributions sorted by Communication Summary.

*Table 79. Report distributions by communication summary*

| Report name        | ID  | Description |
|--------------------|-----|-------------|
| Line elapsed time  | D18 | N/A         |
| Line not wait time | D19 | N/A         |

- Table 80 shows the report distributions sorted by Line Functions.

*Table 80. Report distributions by line functions*

| Report Name              | ID  | Description |
|--------------------------|-----|-------------|
| Received block length    | D36 | N/A         |
| Transmitted block length | D37 | N/A         |
| Inactive intervals       | D38 | N/A         |

- The following table shows the report distributions sorted by MSC Queuing Summary.

*Table 81. Report distributions by MSC queuing summary*

| Report name   | ID  | Description |
|---------------|-----|-------------|
| Time in queue | D39 | N/A         |

- The following table shows the report distributions sorted by Transaction Queuing.

*Table 82. Report distributions by transaction queuing*

| Report name                        | ID  | Description   |
|------------------------------------|-----|---------------|
| Transactions on queue at schedule  | D17 | N/A           |
| Transactions dequeued per schedule | D16 | N/A           |
| Prefetch format blocks             | D28 | Elapsed time  |
|                                    | D29 | Not wait time |

- The following table shows the report distributions sorted by Call Summary.

*Table 83. Report distributions by call summary queuing*

| Report name                           | ID  | Description  |
|---------------------------------------|-----|--------------|
| PSB waits per DL/I call               | D13 | N/A          |
| PSB waits per external subsystem call | D44 | N/A          |
| PSB elapsed time per call             | D11 | N/A          |
| PSB not wait time per call            | D12 | N/A          |
| PSB external subsystem calls          | D45 | Elapsed time |

- The following table lists some distributions derived from buffer pool statistics for wait times.

*Table 84. Wait time distributions*

| Function                 | ID  | Module key |
|--------------------------|-----|------------|
| Storage                  | D22 | SMN        |
| Scheduler internal       | D25 | MSC        |
| Queue manager I/O        | D26 | QMG        |
| Block loader I/O         | D27 | BLR        |
| MFS block I/O            | D32 | MFS        |
| MFS directory I/O        | D33 | MFS        |
| Format buffer pool space | D35 | PMM        |
| QMGR SNAPQ check         | D42 | None       |

## How to get a frequency distribution output

To request the IMS Monitor Report Print utility to gather distribution data, include a DIS input control statement. This causes all report items with an entry under a column headed MEAN to have a corresponding frequency distribution as part of the Distribution Appendix report. Each report line includes an identifying reference number under the column headed Distribution Number. You can use the reference number to locate the distribution data flagged by that number in the appendix.

## How frequency distribution ranges are defined

A set of ten intervals is defined for each summary line and the occurrences falling in each interval are cumulated. The interval ranges are preset with default end points. For example, the end points, for DL/I call elapsed time are: 0, 1000, 2000, 4000, 8000, 16000, 32000, 64000, 128000, 256000, INF (all times are in milliseconds). The default end points are chosen so that they are suitable to the event. The lower limit of the first interval always defaults to zero, and the upper limit of the tenth interval is infinity (INF).

Although several types of distribution can use the same set of end points, each type is assigned a distribution identifier. You can use this to redefine the end points. To override the default end points you include an input control statement to the Report Print utility. The statement specifies the type of distribution identifier and gives the desired end point values.

The DL/I call elapsed time end points could be respecified by:

```
D5  0,500,1000,1500,2000,4000,,,100000,500000
```

The values of the unspecified end points remains at their default values of 32000 and 64000 as does the last (INF).

The following example shows a sample page from the Distribution Appendix report, which gives an example of how ranges vary with the type of distribution. The lines are arranged in pairs, with the second one recording the cumulated counts.

```
IMS MONITOR  ****DISTRIBUTION APPENDIX****  TRACE START 1993 130  5:55:15  TRACE STOP 1993 130  5:59:49 PAGE 0200

#  1.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          1
#  2.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          1
#  3.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          1
#  4.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          1
#  5.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          6
#  6.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          8
#  7.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          1
#  8.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          9
#  9.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          1
# 10.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          1          0          0          0          0          0          0          0          0          0          7
# 11.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          1
# 12.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          8
# 13.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
#          0          0          0          0          0          0          0          0          0          0          1
# 14.....0.....200000.....400000.....600000.....800000.....1000000.....1200000.....1400000.....1600000.....1800000.....INF
```

|   |         |        |             |             |             |             |              |              |              |              |              |     |
|---|---------|--------|-------------|-------------|-------------|-------------|--------------|--------------|--------------|--------------|--------------|-----|
|   |         | 0      | 0           | 0           | 0           | 0           | 0            | 0            | 0            | 0            | 1            |     |
| # | 15..... | 0..... | 200000..... | 400000..... | 600000..... | 800000..... | 1000000..... | 1200000..... | 1400000..... | 1600000..... | 1800000..... | INF |
|   |         | 1      | 0           | 0           | 0           | 0           | 0            | 0            | 0            | 0            | 0            | 8   |
| # | 16..... | 0..... | 200000..... | 400000..... | 600000..... | 800000..... | 1000000..... | 1200000..... | 1400000..... | 1600000..... | 1800000..... | INF |
|   |         | 0      | 0           | 0           | 0           | 0           | 0            | 0            | 0            | 0            | 0            | 1   |

## Default values of distribution definitions

Using an identifier provided in the frequency distribution tables and the wait time distributions table you can determine the default end points for the distribution by locating it in the following list:

**D1, D2, D5, D6, D9, D10, D11, D12, D15 D18, D19, D20, D21, D22, D25, D27, D28, D29, D30, D31, D43, and D45**

0, 1000, 2000, 4000, 8000, 16000, 32000, 64000, 128000, 256000, INF

**D3** 0, 50000, 100000, 150000, 200000, 250000, 300000, 350000, 400000, 450000, INF

**D4** 0, 200000, 400000, 600000, 800000, 1000000, 1200000, 1400000, 1600000, 1800000, INF

**D7, D13, D44**

0, 0, 1, 2, 3, 4, 5, 6, 7, 8, INF

**D8** 0, 100000, 200000, 300000, 400000, 500000, 600000, 700000, 800000, 900000, INF

**D14, D16, D17**

0, 1, 2, 3, 4, 5, 10, 15, 30, 90, INF

**D23, D24, D26, D32, D40, D42**

0, 2000, 8000, 24000, 50000, 100000, 150000, 200000, 250000, 300000, INF

**D33, D34, D35**

0, 2000, 4000, 8000, 16000, 32000, 64000, 96000, 128000, 160000, INF

**D36, D37**

0, 10, 20, 40, 80, 100, 200, 400, 800, 1000, INF

**D38** 0, 1000, 10000, 100000, 200000, 500000, 800000, 1000000, 1500000, 2000000, INF

**D39** 0, 1000, 5000, 10000, 50000, 100000, 500000, 1000000, 5000000, 10000000, INF

## Interpreting Distribution Appendix output

You can use the detailed output in the Distribution Appendix when you suspect an unusual combination of events was reported in a report summary line. Usually, the average and maximum times or counts are sufficient to highlight a resource usage problem.

However, if you suspect the mean value to be masking an unusual distribution you can draw on the detail contained in the IMS Monitor output records.

For example, suppose you are investigating a change in the scheduling algorithm for a particular transaction and need to know how many transactions were able to be processed for each scheduling of an application program. The following figure shows the processed transactions in a histogram:






Figure 94. Number of transactions processed for each scheduling of an application program

The average is 2.5 transactions per schedule. The distribution in the figure suggests that many schedules were able to process only one or two transactions, and few schedules significantly exhausted the queue. The distribution data for the histogram is as follows:

|                       |   |   |   |   |   |      |     |
|-----------------------|---|---|---|---|---|------|-----|
| Number of schedules   | 1 | 2 | 3 | 4 | 5 | 6-10 | >10 |
| Transactions dequeued | 8 | 7 | 0 | 0 | 2 | 1    | 0   |

The Distribution Appendix presents the histogram data in the form of two lines:

- The first line shows the intervals, prefixed by a cross reference to an individual line on the earlier output.
- The second line gives the number of events occurring in those intervals.

This data appears as follows:

```
# 950B...0...1...2...3...4...5...10...15...30...90...INF
      8  7  0  0  2  1  0  0  0  0
```

The cross reference 950B points to a unique report line. For example, the Transaction Queuing report on the appropriate line for the transaction of interest shows 950A,B under the column headed DISTRIBUTION NUMBER. Use the reference number 950B to locate the data in the Distribution Appendix. The 950A reference points to the data for the number of transactions in the queue at schedule time.

---

## Chapter 51. //DFSSTAT reports

The //DFSSTAT reports show you how many DB and DC calls are issued by an application program and describe buffering activity during the application's execution. The reports are written when the application program terminates.

To get the reports, you must put a //DFSSTAT DD statement in the JCL of your batch region or online dependent region. The following is an example of a //DFSSTAT DD statement.

```
//DFSSTAT DD    SYSOUT=A
```

**Recommendation:** Although it is supported, do not include the //DFSSTAT DD statement in the JCL for an MPP region. If you do not include the //DFSSTAT DD statement, you avoid the overhead and large amount of output that results from creating one set of reports each time a short-running MPP terminates.

The following topics describe the reports that //DFSSTAT creates.

### Related concepts:

“Tools for detailed monitoring” on page 378

---

## PST-Accounting report

This report shows how many DB and DC calls are issued by an application program.

### Fields in the PST-Accounting report

The following example shows a PST-Accounting report and shows the names of each field. Each field in this report represents one type of DB or DC call. For example, the DB GU CALLS field shows how many database Get Unique calls were issued by the application.

```
*** PST ACCOUNTING STATISTICS ***
DB GU CALLS                2
DB GN CALLS                2
DB GNP CALLS               1
DB GHU CALLS               1
DB GHN CALLS               1
DB GHNP CALLS              1
DB ISRT CALLS              2
DB DLET CALLS              1
DB REPL CALLS              1
DB CALLS (TOTAL)          12
DB DEQ CALLS               1
DB RLSE CALLS              0
MSG GU CALLS               2
MSG GN CALLS               2
MSG CHNG CALLS             4
MSG ISRT CALLS             8
MSG PURGE CALLS            4
MSG CMD CALLS              1
MSG GCMD CALLS             1
MSG AUTH CALLS             1
MSG SETO CALLS             4
SYS APSB CALLS             0
SYS DPSB CALLS             0
SYS GMSG CALLS             2
```

|                |   |
|----------------|---|
| SYS ICMD CALLS | 1 |
| SYS RCMD CALLS | 2 |
| SYS CHKP CALLS | 0 |
| SYS XRST CALLS | 0 |
| SYS ROLB CALLS | 1 |
| SYS ROLS CALLS | 2 |
| SYS SETS CALLS | 1 |
| SYS SETU CALLS | 1 |
| SYS INIT CALLS | 1 |
| SYS INQY CALLS | 3 |
| SYS LOG CALLS  | 1 |

---

## VSAM-Buffer-Pool report (for batch regions only)

The VSAM-Buffer-Pool report describes VSAM buffer pool activity during the execution of an application program. A separate report is written for each VSAM subpool. The last VSAM-Buffer-Pool report summarizes the buffering activity in all the VSAM subpools used by the application.

### Using the VSAM-Buffer-Pool report

This report is written only for applications that you run in batch regions.

The primary use of the VSAM-Buffer-Pool report, an example is shown in the following example, is to see how many I/O operations were issued in each VSAM subpool.

```

*** VSAM BUFFER POOL STATISTICS ***
FIX INDEX/BLOCK/DATA                N/Y/N
SHARED RESOURCE POOL ID              VPL1
SHARED RESOURCE POOL TYPE            D
SUBPOOL BUFFER SIZE                  4,096
TOTAL BUFFERS IN SUBPOOL              1,000
TOTAL HIPERSPACE BUFFERS IN SUBPOOL   50

RETRIEVE BY RBA CALLS                 370
RETRIEVE BY KEY CALLS                 187583
LOGICAL RECORDS INSERTED INTO ESDS    310
LOGICAL RECORDS INSERTED INTO KSDS    9823
LOGICAL RECORDS ALTERED IN THIS SUBPOOL 0
TIMES BACKGROUND WRITE FUNCTION INVOKED 0
SYNCHRONIZATION CALLS RECEIVED        29923
PERM WRT ERROR BUFFS NOW IN THE SUBPOOL 0
LARGEST NBR OF PERM ERR BUFFS EVEN IN THE SUBPL 0
VSAM GET CALLS ISSUED                  0
VSAM SCHBFR CALLS ISSUED               189290
CONTROL INTERVAL REQUESTED ALREADY IN POOL 0
*CONTROL INTERVAL READ FROM EXTERNAL STORAGE 51238
*VSAM WRITES INITIATED BY IMS          138637
*VSAM WRITES TO MAKE SPACE IN THE POOL 9288
VSAM READS FROM HIPERSPACE BUFFERS    0
VSAM WRITES FROM HIPERSPACE BUFFERS   0
FAILED VSAM READS FROM HIPERSPACE BUFFERS 0
FAILED VSAM WRITES TO HIPERSPACE BUFFERS 0

*TOTAL I/O OPERATIONS                 199163

```

### Fields in the VSAM-Buffer-Pool report

The VSAM-Buffer-Pool report is identical to the VSAM-Buffer-Pool report written by the DB monitor, with the following exceptions:

- Although the field names in the DB Monitor's VSAM Buffer Pool report are preceded by "NUMBER OF", the fields in both reports have the same meaning.
- The //DFSSTAT VSAM-Buffer-Pool report does not keep track of the start trace and end trace times. This is unnecessary because information is always gathered for the //DFSSTAT reports from the beginning to the ending of the application's execution.
- The //DFSSTAT VSAM-Buffer-Pool report contains a "TOTAL I/O OPERATIONS" field, which is the sum of the following:
  - The number of times a CI was read into the buffer from the database (CONTROL INTERVAL READ FROM EXTERNAL STORAGE field in the report)
  - The number of times a buffer was written to the database (VSAM WRITES INITIATED BY IMS field in the report)
  - The number of times a buffer was written to the database so a new CI could be read into the buffer (VSAM WRITES TO MAKE SPACE IN THE POOL field in the report).
  - The //DFSSTAT VSAM-Buffer-Pool report includes a summary report. The summary report is preceded by SUBPOOL BUFFER SIZE=ALL. It contains a summary of read and write information for all VSAM Buffer Pool reports.

The fields that represent I/O operations are highlighted on the left by an asterisk (\*).

**Related reading:** See "VSAM-Buffer-Pool report" on page 709 for a description of the various fields in the report.

---

## OSAM-Buffer-Pool report (for batch regions only)

The OSAM-Buffer-Pool report describes the OSAM buffer pool activity during an application's execution. This report is only written for programs running in batch regions.

### Using the OSAM-Buffer-Pool report

The primary use of the OSAM-Buffer-Pool report, an example is shown in the following example, is to see how many OSAM I/O operations were issued. This report does not, however, show Sequential Buffering (SB) related information.

```

*** OSAM DATA BASE BUFFER POOL STATISTICS ***
FIX BLOCK DATA                                Y/Y
SUBPOOL ID                                    004K
SUBPOOL BUFFER SIZE                           4096
TOTAL BUFFERS IN SUBPOOL                       1000

LOCATE-TYPE CALLS                             1765296
REQUESTS TO CREATE NEW BLOCKS                   0
BUFFER ALTER CALLS                             340800
PURGE CALLS                                    39371
LOCATE-TYPE CALLS, DATA ALREADY IN SUBPOOL    1370897
BUFFERS SEARCHED BY ALL LOCATE-TYPE CALLS     1987604
*READ I/O REQUESTS                             375355
*SINGLE BLOCK WRITES BY BUFFER STEAL RTN        0
*BLOCKS WRITTEN BY PURGE                       150284
TOTAL NBR OF I/O ERRORS FOR THIS SUBPOOL       0
BUFFERS LOCKED DUE TO WRITE ERRORS             0
LOCATE CALLS WAITED DUE TO BUSY ID             1431
LOCATE CALLS WAITED DUE TO BUFR BUSY WRITE     0
LOCATE CALLS WAITED DUE TO BUFR BUSY READ     0

```

|                                            |        |
|--------------------------------------------|--------|
| BUFR STEAL/PURGE WAITED FOR OWNERSHIP RLSE | 296    |
| BUFFER STEAL REQUEST WAITED FOR BUFFERS    | 0      |
| *TOTAL I/O OPERATIONS                      | 525639 |

## Fields in the OSAM-Buffer-Pool report

The figure is an example of an OSAM-Buffer-Pool report.

The OSAM-Buffer-Pool report is identical to the Database-Buffer-Pool report written by the DB monitor, with the following exceptions:

- The field names in the IMS Monitor's Database-Buffer-Pool report are preceded by "NUMBER OF", although the fields have the same meaning in both reports.
- The //DFSSTAT OSAM-Buffer-Pool report does not keep track of the start trace and end trace times. This is unnecessary because information is always gathered for the //DFSSTAT reports from the beginning to the ending of the application's execution.
- In addition, the //DFSSTAT OSAM-Buffer-Pool report contains a "TOTAL I/O OPERATIONS" field, which equals the sum of the following fields:
  - READ REQUESTS ISSUED
  - OSAM WRITES ISSUED
  - QUEUED WRITES ISSUED
  - FORMAT LOGICAL CYLINDER REQUESTS
  - BISAM READS OR QISAM SETLS.

These fields represent I/O operations and are highlighted on the left by an asterisk (\*).

---

## Sequential-Buffering-Summary report

The Sequential-Buffering-Summary report provides an overview of SB-related information for the application. (VSAM-related information is not included.)

### Using the Sequential-Buffering-Summary report

From an SB Summary report, an example is shown in the following example, you can determine if the application benefited from the use of SB. When using this report, pay particular attention to these fields:

- NBR BLOCKS READ SEQUENTIALLY and PCT OF TOTAL
- PERCENT READ PER SEARCH REQUEST

\*\*\* SEQUENTIAL BUFFERING SUMMARY FOR THE APPLICATION \*\*\*

|                                                      |     |
|------------------------------------------------------|-----|
| DFSSBUX0 DISALLOWED USAGE OF SB:                     | NO  |
| DFSSBUX0 REQUESTED CONDITIONAL SB ACTIVATION:        | NO  |
| AT LEAST ONE SB= KEYWORD IN PSB:                     | YES |
| AT LEAST ONE SBPARM CONTROL STMT FOR APPLICATION:    | NO  |
| SBPARM CONTROL CARD(S) READ FROM //DFSCTL:           | YES |
| AT LEAST ONE SBPARM PSB= SPECIFIED THAT MATCHED PSB: | YES |
| AT LEAST ONE SBPARM DB= SPECIFIED THAT MATCHED DB:   | YES |
| AT LEAST ONE SBPARM PCB= SPECIFIED THAT MATCHED PCB: | NO  |
| AT LEAST ONE SBPARM DD= SPECIFIED THAT MATCHED DD:   | NO  |

|                                              |       |
|----------------------------------------------|-------|
| NUMBER OF SEARCH REQUESTS ISSUED BY OSAM BH: |       |
| SEARCH                                       | 2,213 |
| NUMBER OF READ I/O:                          |       |
| RANDOM READ                                  | 686   |
| SEQUENTIAL READ                              | 652   |

|                                 |       |                 |
|---------------------------------|-------|-----------------|
| NUMBER OF BLOCKS READ:          |       |                 |
| TOTAL NUMBER BLOCKS READ        | 7,206 |                 |
| NBR BLOCKS READ AT RANDOM       | 686   | PCT OF TOTAL:   |
| NBR BLOCKS READ SEQUENTIALLY    | 6,520 | PCT OF TOTAL: 9 |
| PERCENT READ PER SEARCH REQUEST | 60.46 |                 |
| NUMBER OF SEQUENTIAL I/O ERRORS | 0     |                 |

## Fields in the Sequential-Buffering Summary report

The example shows a Sequential-Buffering-Summary report.

The first part of this report shows why Sequential Buffering was or was not used. This part of the report describes whether:

- A SBONLINE control card was provided in DFSVSMxx (this applies only to IMS DC environments).
- A /STOP SB command was in effect when the application program started (this applies only to IMS DC environments).
- The SB Initialization Exit Routine (DFSSBUX0) disallowed use of SB.
- The SB Initialization Exit Routine (DFSSBUX0) requested conditional activation of SB by default.
- At least one SB= keyword was provided during PSBGEN.
- The //DFSCTL file contained at least one SBPARM control statement that applied to the application program.
- SBPARM control cards have been read. If the answer is Yes, the following statistics indicate what SBPARM keywords were used. This can be helpful in determining why sequential buffering was or was not used for the application program.
- At least one PSB= keyword was specified on an SBPARM control card and it matched the PSB used by the application.
- At least one DB= keyword was specified on an SBPARM control card where the PSB matched or was not specified, and the database matched one used by the application.
- At least one PCB= keyword was specified on an SBPARM control card where the PSB and DB matched or were not specified, and the PCB name matched one used by the application.
- At least one DD= keyword was specified on an SBPARM control card where the PSB, DB, and PSB matched or were not specified, and the DD name matched one used by the application.
- Whether SBPARM control cards have been read. If the answer is "yes," the following statistics indicate what SBPARM keywords were used. This is helpful in determining why sequential buffering was or was not used for the application program.
- At least one PSB= keyword was specified on a SBPARM control card and it matched the PSB used by the application.

The other fields in the report are as follows:

### NUMBER OF SEARCH REQUESTS ISSUED BY THE OSAM BH

This field shows you how many times the OSAM buffer handler asked the SB buffer handler to search the SB buffer pools for a specific OSAM block.

The value in this field is equal to the number of OSAM random read I/O operations that would have been issued without SB.

### NUMBER OF READ I/O

These fields show you the number of OSAM random and sequential read I/O operations it took to satisfy requests made by the application program. The sum of these two numbers is the total number of OSAM read I/O operations issued on behalf of the application. You can subtract this sum from the NUMBER OF SEARCH REQUESTS ISSUED BY THE OSAM BH field to calculate how many read I/O operations you saved by using SB.

### NUMBER OF BLOCKS READ

These fields tell you how many OSAM data set blocks were read to satisfy requests from the application program. These fields show you:

- The total number of blocks read
- The number and percentage of blocks read with a random read
- The number and percentage of blocks read with a sequential read

If the percentage of blocks read with a sequential read is high, SB probably helped reduce the elapsed time of the application program.

### PERCENT READ PER SEARCH REQUEST

This field shows you the number of read I/O operations issued by the SB buffer handler expressed as a percentage of the number of times the OSAM buffer handler asked the SB buffer handler to search for a block.

A low percentage indicates that many of the search requests were satisfied without issuing an I/O operation. Therefore, a low number in this field shows that SB probably helped reduce the elapsed time of the application program.

### NUMBER OF SEQUENTIAL I/O ERRORS

This field describes the number of sequential reads that resulted in I/O errors. When an I/O error is detected during a sequential read, IMS increments this field and marks the 10 SB buffers involved in the read as invalid. Then IMS issues a random read for the block that was requested by the OSAM buffer handler.

---

## Sequential Buffering Detail report

This report gives you detailed information about how SB was used for a particular SB buffer pool. A separate report is created for each SB buffer pool used by the application program.

Each report consists of three pages, A, B, and C. Summary information is contained on page A. More detailed information is found on pages B and C.

Each of the following topics contains an example of the relevant page (A, B or C) of the Sequential-Buffering-Detail report.

### Fields on page A of the Sequential Buffering Detail report

The following example shows page A of the report.

```
//DFSSTAT STATISTICS FOR: JOB=OSBTC01 STEP=STEP1 . PGM=DFSDDLTO PSB=PBVDSALR DATE=93.058 TIME=09.39
```

```
-----  
*** SB DETAIL STATISTICS (PAGE A) ***  
PSB          PBVDSALR  
DB            DBOVLFP  
PCB  
DB-PCB NBR      1  
DSG-CB NBR      1  
DD            VLOSAM01  
DB-ORG          HDAM  
DD-TYPE        *PSDATA  
NBR OF BUFSETS      4
```

```

COMPARE-OPTION IS ACTIVE
** NUMBER OF SEARCH REQUESTS ISSUED BY OSAM BH:
    SEARCH                                2,213
** NUMBER OF READ I/O:
    TOTAL                                1,338
    RANDOM READ                          686
    SYNCHRONOUS SEQUENTIAL READ          555
    OVERLAPPED SEQUENTIAL READ           97
** NUMBER OF BLOCKS READ:
    TOTAL                                7,206
    RANDOM READ                          686
    SYNCHRONOUS SEQUENTIAL READ          5,550
    OVERLAPPED SEQUENTIAL READ           970
    PCT OF TOTAL:    9.51
    PCT OF TOTAL:   77.01
    PCT OF TOTAL:   13.46
** AVERAGE I/O WAIT TIMES (MILLIS):
    RANDOM READ                          15.70
    SYNCHRONOUS SEQUENTIAL READ          18.03
    OVERLAPPED SEQUENTIAL READ           .26

```

You can use the first topic of this page to identify the SB buffer pool, database PCB, and DD name this report applies to. There can be more than one SB buffer pool (and, thus, more than one report) created for a particular database PCB and DD name. This can happen if the PCB is involved in a logical relationship.

This topic contains the following information:

- The PSB name.
- The DBD name (as coded in the PCB macro during PSBGEN).
- The PCB label (as coded in the PCB macro during PSBGEN).
- The relative number of the database PCB within the PSB.
- A unique identifier of the data set group control block within the database PCB. (You can use this number to uniquely identify the SB buffer pool when more than one SB buffer pool has been created for the same PCB and DD name.)
- The DD name.
- The type of database organization (HDAM, HIDAM, PHDAM, PHIDAM, or HISAM).
- The type of database data set. This can be one of three values:
  - \*INDX indicates this report applies to a data set used as an index.
  - \*PSDATA indicates this report applies to a data set containing data accessed according to its primary sequence.
  - \*SSDATA indicates this report applies to a data set containing data accessed according to a secondary index sequence or by crossing a logical relationship.
- The number of buffer sets in the SB buffer pool. The default number of buffer sets is four. You might have changed this default, however, with a SBPARM control statement or in the SB Initialization Exit routine.

#### **NUMBER OF SEARCH REQUESTS ISSUED BY OSAM BH**

This field shows you the number of OSAM random read I/O operations that would have been issued without using SB.

#### **NUMBER OF READ I/O**

These fields show you the number of read I/O operations that were actually issued in this SB buffer pool. The fields tell you:

- The total number of OSAM read I/O operations
- The number of random reads
- The number of synchronous sequential reads
- The number of overlapped (asynchronous) sequential reads

You can use these fields to calculate the percentage of each of the three types of read I/O operations that were used for this database PCB and DD name.



You can also subtract the total number of OSAM read I/O operations from the NUMBER OF SEARCH REQUESTS ISSUED BY OSAM BH field to determine the number of read I/O operations you saved by using SB.

#### NUMBER OF BLOCKS READ

These fields show you how many blocks were read by each type of read I/O operation. The fields tell you:

- The total number of blocks read
- The number and percentage of blocks read by random reads
- The number and percentage of blocks read by synchronous sequential reads
- The number and percentage of blocks read by overlapped (asynchronous) sequential reads

A high percentage of blocks read with sequential reads indicates that SB probably helped the application program run faster, at least while processing this database PCB and data set.

A high percentage of blocks read with random reads, however, might indicate:

- A large amount of random processing by the application program
- The OSAM data set associated with this database PCB needs reorganizing

If many blocks were read with random reads, and if the program was processing the database sequentially, you can sometimes get better buffering performance by increasing the number of buffer sets (for example, increase the value of the BUFSETS parameter on the SBPARM control statement in //DFSCTL). After increasing the number of buffer sets, observe how the number reported in the NUMBER OF READ I/O field changes the next time the application is executed.

A small percentage of blocks read at random sometimes indicates that you can reduce the number of buffer sets to save virtual storage space.

#### AVERAGE I/O WAIT TIMES (MILLIS)

These fields tell you the average I/O wait times for each of the three types of read I/O operations. These times show the average time the application program waited for a read I/O operation to complete before it could process the data being read.

The times in these fields are in milliseconds. A millisecond is one thousandth of a second (in other words, 50 milliseconds equals 0.050 seconds).

These fields show you:

- The average I/O wait time for random reads
- The average I/O wait time for synchronous sequential reads
- The average I/O wait time for overlapped (asynchronous) sequential reads

These times are only measured when SB is active or IMS is monitoring the I/O reference pattern.

### Fields on page B of the Sequential Buffering Detail report

The following example shows page B of the report.

```
-----
*** SB DETAIL STATISTICS: REFERENCE STATISTICS (PAGE B) ***
** REFERENCES IN BUFFER-SETS:
    RATIO                                .23
** REFERENCES IN RANDOM SRAN CBS:
    RATIO                                .20
** RANDOM SRAN CBS WHICH HAVE BEEN CONVERTED:
    NUMBER                               0
```

| ***** DISTRIBUTION OF REFERENCES IN BUFFER-SETS ***** |                    |                    |              |  |
|-------------------------------------------------------|--------------------|--------------------|--------------|--|
| REFERENCE COUNT                                       | NBR OF OCCURRENCES | PCT OF OCCURRENCES | ACCUMUL. PCT |  |
| 0                                                     | 3                  | .46                | .46          |  |
| 1                                                     | 551                | 84.50              | 84.96        |  |
| 2                                                     | 0                  | .00                | 84.96        |  |
| 3                                                     | 0                  | .00                | 84.96        |  |
| 4                                                     | 0                  | .00                | 84.96        |  |
| 5                                                     | 0                  | .00                | 84.96        |  |
| 6                                                     | 0                  | .00                | 84.96        |  |
| 7                                                     | 0                  | .00                | 84.96        |  |
| 8                                                     | 0                  | .00                | 84.96        |  |
| 9                                                     | 4                  | .61                | 85.58        |  |
| => 10                                                 | 94                 | 14.41              | 100.00       |  |

| ***** DISTRIBUTION OF REFERENCES IN RANDOM SRAN CBS ***** |                    |                    |              |  |
|-----------------------------------------------------------|--------------------|--------------------|--------------|--|
| REFERENCE COUNT                                           | NBR OF OCCURRENCES | PCT OF OCCURRENCES | ACCUMUL. PCT |  |
| 0                                                         | 0                  | .00                | .00          |  |
| 1                                                         | 0                  | .00                | .00          |  |
| 2                                                         | 1                  | 100.00             | 100.00       |  |

This page can be used to evaluate the efficiency of the SB algorithm that monitored the I/O reference pattern for this SB buffer pool. By analyzing this page, you can answer these questions:

- How efficient were the decisions to issue sequential reads? This question can be answered by the REFERENCES IN BUFFER SETS field.
- How efficient were the decisions to issue random reads? This question can be answered by the REFERENCES IN RANDOM SRAN CBS field and the RANDOM SRAN CBS WHICH HAVE BEEN CONVERTED field.

The fields on Page B reflect buffering activity when SB was active. Buffering activity when IMS was only monitoring the I/O reference pattern is not included.

#### REFERENCES IN BUFFER SETS

This field shows you how many times blocks read with a sequential read were referenced by the OSAM buffer handler. This ratio is calculated as follows:

The number of references in buffer sets divided by the number of blocks read by sequential read.

For example, a ratio of 1.00 means that, on the average, each block read by a sequential read was referenced once. A ratio of 0.50 means that, on the average, only half the blocks read by a sequential read were referenced once.

In general, a high ratio (for example, 0.85) shows you that the decisions to issue sequential reads were efficient and probably helped reduce the execution time of the application. A low ratio (for example, 0.30) shows you that the benefit of issuing sequential reads was smaller because many of the blocks read were never referenced.

#### REFERENCES IN RANDOM SRAN CBS

This field shows you the effectiveness of the decisions to issue random reads. To understand what this ratio means, you must first understand how the SB buffer handler uses control blocks to track the I/O reference pattern.

The SB buffer handler uses a control block called SB range (SRAN) to track the I/O reference pattern within a range of 10 consecutive blocks. Each SRAN has a reference counter that shows how many times the OSAM buffer handler requested a block contained in the set of 10 blocks tracked by the SRAN.

There are two types of SRAN control blocks. One type, called a *sequential* SRAN, maintains counts for 10 consecutive blocks that have been read with a sequential read. A second type, called a *random* SRAN, maintains counts for 10 consecutive blocks that have been referenced in a random pattern. There is one sequential SRAN for each buffer set in an SB buffer pool; the number of random SRANs is twice the number of sequential SRANs.

The SRANs are chained together on a “use chain,” that is, they are chained together in the order in which they have been most or least recently used. Each type of SRAN has its own use chain. Most recently used SRANs are at the top of the use chain; least recently used SRANs are at the bottom. Whenever a SRAN is needed to track a set of blocks that are not currently being tracked, the SB buffer handler selects a SRAN from the bottom of the use chain. This is because ranges of consecutive blocks that have not been referenced recently are less likely to be referenced again in the near future.

Each time a random SRAN is reused, the reference count maintained in the SRAN is recorded for later use. When the application program terminates, these reference counts are used to calculate the number in this field. The number is calculated as follows:

$$\text{Ratio} = X / (Y * 10)$$

where,

**X**        Total number of random references in random SRANs (this number is usually equal or close to the number of random reads in the sets of blocks tracked by random SRANs)

**Y**        Total number of reused random SRANs

A low ratio (for example, 0.15) in this field means that on the average, the SB buffer handler correctly recognized random I/O reference patterns. A ratio of 0.15 means that the SB buffer handler issued an average of 1.5 random reads for blocks tracked by a random SRAN. Issuing 1.5 random reads for a set of 10 consecutive blocks is normally more efficient than issuing one sequential read for the 10 consecutive blocks.

A high ratio (for example, 0.50) in this field probably indicates that the SB buffer handler often mistook sequential reference patterns as random reference patterns. A ratio of 0.50 means that the SB buffer handler issued an average of 5 random reads for blocks tracked by a random SRAN. Issuing 5 random reads for a set of 10 consecutive blocks is normally less efficient than issuing one sequential read for the 10 consecutive blocks.

#### **RANDOM SRAN CBS WHICH HAVE BEEN CONVERTED**

These fields are another measurement of the SB buffer handler's effectiveness in analyzing the I/O reference patterns.

Sometimes the SB buffer handler interprets a reference to a set of consecutive blocks as a random pattern, then, after some additional references, detects that the set of blocks is actually being referenced sequentially. In this case, the SB buffer handler first issues several random reads in a set of consecutive blocks and later issues a sequential read for the same set of blocks. When this occurs, the random SRAN that tracks the set of blocks is converted to a sequential SRAN. If the SB buffer handler had originally read the set of blocks with a sequential read, the cost of issuing several random reads would have been avoided.

The fields in this topic are:

**NUMBER**

This indicates how many times random SRANs were converted to sequential SRANs during the application program's execution.

**PCT OF STOLEN RANDOM SRAN**

The value in this field expresses the number of converted SRANs as a percentage of the number of times the SB buffer handler had to acquire a random SRAN from the use chain. A high percentage (for example, 40 percent) probably indicates that the SB buffer handler often mistook a sequential reference pattern as random reference pattern.

**DISTRIBUTION OF REFERENCES IN BUFFER SETS**

This table shows you more detailed information about the ratio reported in the REFERENCES IN BUFFER SETS field.

For example, the REFERENCES IN BUFFER SETS field might indicate that 80 percent (0.80) of the blocks in the SB buffer sets were referenced. What does this mean? It could mean that 20 percent of the time none of the blocks in a buffer set were referenced. Or instead, it could mean that 100 percent of the time only 8 out of 10 blocks in the buffer sets were referenced. You can answer this question by analyzing this table.

The fields in this table are:

**REFERENCE COUNT**

This column lists the number of references to blocks in a buffer set. A zero in this column means that no references were made to blocks in a buffer set, a one means that one reference was made, and so on. The "=> 10" in the last row of this column means that 10 or more references were made.

**NBR OF OCCURRENCES**

This column shows you how many buffer sets were referenced the number of times shown in the REFERENCE COUNT column. For example, in the figure, the first number in this column is 3, which says that 3 buffer sets were never referenced. The second number says that 551 buffer sets were referenced once.

**PCT OF OCCURRENCES**

This column shows you the value in the NBR OF OCCURRENCES column expressed as a percentage of the number of times buffer sets were reused. For example, in the figure, the first number in this column is 0.46, which says that 0.46 percent of the reused buffer sets were never referenced. The second number says that 84.50 percent of the reused buffer sets were referenced once.

**ACCUMUL. PCT**

This column shows you the accumulated PCT OF OCCURRENCES from zero through the current reference count. For example, in the figure, the third number in this column is 84.96, which says that 84.96 percent of the reused buffer sets were referenced two or less times.

**DISTRIBUTION OF REFERENCES IN RANDOM SRAN CBS**

This table shows you more detailed information about the ratio reported in the REFERENCES IN RANDOM SRAN CBS field. The table is similar to the DISTRIBUTION OF REFERENCES IN BUFFER SETS table, except that it shows information about references tracked by random SRAN control blocks.

**REFERENCE COUNT**

This column lists the number of references to blocks tracked by random SRANs. A zero in this column means that no references were made and, therefore, no random reads were issued for blocks tracked by a random

SRAN. A “one” means that one reference was made to blocks tracked by a random SRAN (on the average, one reference for every tenth block tracked by the random SRAN).

#### **NBR OF OCCURRENCES**

This column shows you how many random SRANs were referenced the number of times shown in the REFERENCE COUNT column.

#### **PCT OF OCCURRENCES**

This column shows you the NBR OF OCCURRENCES expressed as a percentage of the number of times that random SRANs were reused.

#### **ACCUMUL. PCT**

This column shows you the accumulated PCT OF OCCURRENCES from zero through the current reference count.

### **Fields on page C of the Sequential Buffering Detail report**

The following example shows page C of the report.

```
//DFSSTAT STATISTICS FOR: JOB=OSBTC01 STEP=STEP1 . PGM=DFSDDLTO PSB=PBVDSALR DATE=93.058 TIME=09.39
-----
*** SB DETAIL STATISTICS: INTERNAL COUNTERS AND VALUES (PAGE C) ***

** DEACTIVATIONS:
   NBR OF SB-DEACTIVATION                1
   NBR OF MONITORING-DEACTIVATION         0
** RESULTS OF EVALUATION OF SEQUENTIALITY:
   NBR POSITIVE RESULTS                   3
   NBR NEGATIVE RESULTS                   1
** RESULTS OF EVALUATION OF ACTIVITY RATE:
   NBR POSITIVE RESULTS                   4
   NBR NEGATIVE RESULTS                   0
** NBR RANDOM READ:
   DURING SEQUENTIAL BUFFERING PHASES     27
   DURING "MONITORING-ONLY" PHASES       659
   WHILE NOT MONITORING REFERENCE PATTERN  0
** NBR RANDOM READS WITH SEQUENTIAL REFERENCE PATTERN:
   ACCESS TO INVALID BUFFERS              25
   ACCESS AT DATA SET END                 2
** NBR OF BUFFERING POSITIONS:            2,213
** INTERNAL SB-ALGORITHM VALUES:
   SDSGBPTR: BLOCKS PER TRACK              31
   SDSGNBRB: BLOCKS PER BUFSET             10
   SDSGSCST: RELATIVE SEQ I/O COSTS        1.36
   SDSGSINB: SIZE OF NEIGHBORHOOD          2
   SDSGTHR1: THRESHOLD CURRENT+1           2
   SDSGTHR2: THRESHOLD OVERLAP             3
   SDSGTHR3: THRESHOLD NEIGHB             11
```

Page C consists of the following topics:

#### **DEACTIVATIONS**

The fields in this topic are:

##### **NBR OF SB DEACTIVATION**

This field shows how many times SB was deactivated.

##### **NBR OF MONITORING DEACTIVATION**

This field shows how many times I/O reference monitoring was deactivated.

Deactivation of I/O reference monitoring can occur for one of two reasons:

- If several consecutive periodical evaluations of the buffering process show that it is not worthwhile to use SB
- If you set a limit on the SB buffer space by specifying the MAXSB keyword in the SBONLINE control statement. The use of SB is restricted when this limit is reached.

## **RESULTS OF EVALUATION OF SEQUENTIALITY**

This topic and the next topic, RESULTS OF EVALUATION OF ACTIVITY RATE, help explain why SB was deactivated (or was not activated) during the application program's execution.

The decision to activate and deactivate SB is made by a periodical evaluation of the buffering process for a particular DB-PCB/DSG control block pair. The evaluation is based on the following criteria:

- Sequentiality
- Activity rate

If the results for both tests are positive (in other words, if the results of both tests recommend use of SB), IMS will activate SB (if not active) or will continue to use SB. If at least one of the test results is negative, then IMS will deactivate SB (if active) or will continue not to use SB. After several decisions not to use SB, IMS can also deactivate monitoring of the I/O reference pattern.

The fields in this topic are:

### **NBR POSITIVE RESULTS**

This field shows you how many times periodical evaluation of the I/O reference pattern detected enough of a sequential reference pattern to warrant use of SB.

### **NBR NEGATIVE RESULTS**

This field shows you how many times periodical evaluation of the I/O reference pattern did not detect enough of a sequential reference pattern to warrant use of SB.

## **RESULTS OF EVALUATION OF ACTIVITY RATE**

The fields in this topic are:

### **NBR POSITIVE RESULTS**

This shows you how many times periodical evaluation detected an I/O activity rate high enough to warrant use of SB.

### **NBR NEGATIVE RESULTS**

This shows how many times periodical evaluation determined that the I/O activity rate was not high enough to warrant use of SB.

## **NBR RANDOM READ**

This topic shows you how many random reads were issued during each of the following types of buffering phases:

### **DURING SEQUENTIAL BUFFERING PHASES**

This field shows you how many random reads were issued while SB was active.

### **DURING "MONITORING ONLY" PHASES**

This field shows you how many random reads were issued while SB was not active and IMS was still monitoring the I/O reference pattern.

### **WHILE NOT MONITORING REFERENCE PATTERN**

This field shows you how many random reads were issued while SB was not active and IMS was not monitoring the I/O reference pattern.

## **NBR RANDOM READS WITH SEQUENTIAL REFERENCE PATTERN**

This topic shows you how many times random reads were issued even though the I/O reference pattern was sequential. These counters are updated only when SB is active.

The fields in this topic are:

#### **ACCESS TO INVALID BUFFERS**

This field shows how many times a random read was issued because the contents of an SB buffer was invalid and could not be used. Some possible reasons for marking an SB buffer invalid are:

- Your IMS system is running in a block-level sharing environment. In a block-level sharing environment, more than one IMS system can read from and write to the same database. For example, if IMS system “A” reads a block into a buffer and IMS system “B” updates that block while the block is still in system “A's” buffer, system “A's” buffer will be marked as invalid.
- The activation and deactivation of SB during periodical evaluations marks SB buffers as invalid.
- A block was being written by another online application in the same IMS subsystem at the same time the SB buffer handler was reading it.
- The SB buffer was marked invalid because of I/O errors.

#### **ACCESS AT DATA SET END**

This field shows how many times the SB buffer handler issued a random read because the set of 10 blocks containing the referenced block was not completely formatted. SB never issues a sequential read at the end of a data set if the last set of consecutive blocks is not completely formatted.

#### **NBR OF BUFFERING POSITIONS**

This field shows how many times the SB buffer handler assumed the application program issued a DL/I call requesting a new position in the database. For example, most GU calls qualified on the key field of a root segment other than the current root segment will cause this counter to be incremented. This counter is maintained only while SB is active and IMS is monitoring the I/O reference pattern.

A high value in this field can indicate a large amount of logical random processing by the application. If other fields seem to indicate that the application did not benefit from SB, this field can explain why.

#### **INTERNAL SB ALGORITHM VALUES**

This topic shows the values of internal counters. They are included to help IMS development during SB problem determination.



---

## Chapter 52. Statistical-analysis, log-transaction reports, and analyzing log records

IMS provides several utilities that extract data from IMS system logs.

- The Statistical Analysis utility produces summary reports of message activity, reports for lines and terminals, and the Message Select and List Report.
- The Log Transaction Analysis utility gives detailed information of individual transaction and processing activities.
- The Knowledge-Based Log Analysis (KBLA) collection of utilities provide an ISPF user interface to perform log record selection and analysis by invoking internal routines and other IMS utility programs.

---

### Statistical analysis utility reports

The input data for the Statistics Analysis utility is a set of IMS log data sets, or user data sets created by the Log Archive utility.

Each input data set can consist of multiple volumes. Several data sets can be concatenated and you can include data sets from multiple IMS systems connected through shared queues or MSC.

**Restriction:** You cannot use system log output from a batch system.

You can use the Transaction Analysis utility to obtain new system logs with reduced content to save processing time.

The Statistical Analysis utility has six control statements you can use to select a subset of transaction activity.

#### Transaction code control statement

You can use this control statement to select a specific transaction code or groups of transaction codes.

#### Symbolic terminal name control statement

You can use this control statement to specify the LTERM name or a generic name. For example, L3270M selects all messages originating from or directed to that LTERM. A generic name of L3270\* could select L3270M and L3270B messages, the comparison being based on the characters preceding the \*.

You can further qualify the output LTERM so that only messages to a given symbolic name resulting from the input LTERM specified are selected.

#### Time control statement

You can specify an interval as a criterion for selection. You give the start and stop times in the form YYDDD and HHMM (Julian day and clock time in minutes). This range criterion is applied to all messages selected by transaction code and terminal specifications.

#### Message select output order statement

Use this control statement to determine which order to list messages. This control statement affects only the content of the Message Select and List or Copy output.



### Nonprintable character control statement

If you anticipate non-printable characters in the message text, you can specify they be printed in hexadecimal format (first character above the second). Otherwise, the characters appear as blanks.

For more information about the Statistical Analysis utility (DFSISTS0), see *IMS Version 11 System Utilities*.

## Calculating transaction loads

There are two reports produced by the Statistical Analysis utility that summarize the distribution of transaction activity across a 24-hour day. Input and output message distributions are separately tabulated for each transaction code and for each device.

A further report shows the response times for each transaction type expressed as percentiles. The reported data is dependent on the selection of system log data that makes up the utility input. The scope of the report can be further limited by selecting a subset of transactions and line traffic as well as the reporting interval.

The following figure shows the format of the Line and Terminal report. For each device on a line the LTERM name is given and a pair of rows of results for send and receive activity is given. The "Total Messages" column is followed by the total and average size of the messages in bytes. A series of hourly intervals divides the 24-hour day and counts of the transaction active in those intervals are recorded.

| LINE AND |     | TERMINAL |            | REPORT |        | DATE 06/05/07 |       |       |       |              |       |       |       |       |       |       |       |       |  | PAGE 00002 |  |  |  |
|----------|-----|----------|------------|--------|--------|---------------|-------|-------|-------|--------------|-------|-------|-------|-------|-------|-------|-------|-------|--|------------|--|--|--|
| NODE     | R/S | TOTAL    | TOTAL      | AVG    | HOURLY |               |       |       |       | DISTRIBUTION |       |       |       |       |       |       |       |       |  |            |  |  |  |
|          |     | MESSAGES | CHARACTERS | SIZE   | 00-07  | 07-08         | 08-09 | 09-10 | 10-11 | 11-12        | 12-13 | 13-14 | 14-15 | 15-16 | 16-17 | 17-18 | 18-19 | 19-24 |  |            |  |  |  |
| CTRL     | R   | 5        | 782        | 156    | 0      | 0             | 0     | 0     | 0     | 5            | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
|          | S   | 10       | 1,123      | 112    | 0      | 0             | 0     | 0     | 0     | 10           | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| CTRLA01  | R   | 3        | 365        | 121    | 0      | 0             | 0     | 0     | 0     | 3            | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
|          | S   | 2        | 276        | 138    | 0      | 0             | 0     | 0     | 0     | 2            | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| CTRLA02  | R   | 2        | 186        | 93     | 0      | 0             | 0     | 0     | 0     | 2            | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
|          | S   | 4        | 306        | 76     | 0      | 0             | 0     | 0     | 0     | 4            | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| L62MVS1  | R   | 9        | 366        | 40     | 0      | 0             | 0     | 0     | 0     | 9            | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
|          | S   | 9        | 483        | 53     | 0      | 0             | 0     | 0     | 0     | 9            | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| SEGUNDO  | S   | 28       | 1,842      | 65     | 0      | 0             | 0     | 0     | 0     | 28           | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| SYSTEM   | R   | 19       | 1,699      | 89     | 0      | 0             | 0     | 0     | 0     | 19           | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| TOTALS   | S   | 53       | 4,030      | 76     | 0      | 0             | 0     | 0     | 0     | 53           | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |

**Key:** R/S—Received/Sent

Entries for devices restricted to input only or output only traffic show only one reporting line. The following example shows the format of the Transaction report. This organizes the data like the Line and Terminal report, except that it is ordered by transaction code.

| TRANSACTION REPORT |       |          |            |        |       |       |       |       |       | DATE 06/05/07 |       |       |       |       |       |       |       |       |  | PAGE 00003 |  |  |  |
|--------------------|-------|----------|------------|--------|-------|-------|-------|-------|-------|---------------|-------|-------|-------|-------|-------|-------|-------|-------|--|------------|--|--|--|
| TRANSACTION        | TOTAL | TOTAL    | AVG        | HOURLY |       |       |       |       |       | DISTRIBUTION  |       |       |       |       |       |       |       |       |  |            |  |  |  |
| CODE               | R/S   | MESSAGES | CHARACTERS | SIZE   | 00-07 | 07-08 | 08-09 | 09-10 | 10-11 | 11-12         | 12-13 | 13-14 | 14-15 | 15-16 | 16-17 | 17-18 | 18-19 | 19-24 |  |            |  |  |  |
| CONV21V0           | R     | 5        | 246        | 49     | 0     | 0     | 0     | 0     | 0     | 5             | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
|                    | S     | 4        | 298        | 74     | 0     | 0     | 0     | 0     | 0     | 4             | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| SMQR21C0           | R     | 3        | 90         | 30     | 0     | 0     | 0     | 0     | 0     | 3             | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
|                    | S     | 2        | 50         | 25     | 0     | 0     | 0     | 0     | 0     | 2             | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| TRAN21V0           | R     | 1        | 30         | 30     | 0     | 0     | 0     | 0     | 0     | 1             | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
|                    | S     | 1        | 25         | 25     | 0     | 0     | 0     | 0     | 0     | 1             | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| SYSTEM             | R     | 9        | 366        | 40     | 0     | 0     | 0     | 0     | 0     | 9             | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |
| TOTALS             | S     | 7        | 373        | 53     | 0     | 0     | 0     | 0     | 0     | 7             | 0     | 0     | 0     | 0     | 0     | 0     | 0     | 0     |  |            |  |  |  |

The following example shows the format of the Transaction report. This report gives the longest and shortest response times for each transaction code for the data selected as input from a set of system logs. Four columns record the percentile response times in seconds. The 25th, 50th, 75th, and 95th percentiles are given. For example, a response time within the 50th percentile is greater than or equal to 50% of the total number of response times processed for that transaction. The first line of response times given is from the completion of the receipt of the input message

until the response message is successfully dequeued. In the event that an output message takes a significantly long time to be completely received at the terminal, a second line shows the receipt to the time the response message is started.

| TRANSACTION      |                 | RESPONSE REPORT  |              |              |              | DATE 06/05/07 |                   | PAGE 00004 |  |
|------------------|-----------------|------------------|--------------|--------------|--------------|---------------|-------------------|------------|--|
| TRANSACTION CODE | TOTAL RESPONSES | LONGEST RESPONSE | 95% RESPONSE | 75% RESPONSE | 50% RESPONSE | 25% RESPONSE  | SHORTEST RESPONSE |            |  |
| CONV21V0         | 4               | .36S             | .31S         | .31S         | .24S         | .19S          | .19S              |            |  |
|                  | 4               | .36S             | .31S         | .31S         | .24S         | .19S          | .19S              |            |  |
| SMQR21C0         | 2               | .51S             | .33S         | .33S         | .33S         | .33S          | .33S              |            |  |
|                  | 2               | .51S             | .33S         | .33S         | .33S         | .33S          | .33S              |            |  |
| TRAN21V0         | 1               | .36S             | .36S         | .36S         | .36S         | .36S          | .36S              |            |  |
|                  | 1               | .36S             | .36S         | .36S         | .36S         | .36S          | .36S              |            |  |

## Assessing program-to-program traffic

When a message processing program directs an output message to another program, that secondary transaction is queued. The transaction code is sometimes unique for convenience of the processing program's logic.

Otherwise, the secondary transaction is queued along with any messages from terminal origin.

You can use the two Messages—Program-to-Program reports to separately count transaction traffic. The following example illustrates the two tabulations. The column headed “Destination” appears above a list of transaction codes that were queued to another program. The originating program is not identified. The column headed “Transaction Code” appears above a list of the initial transaction codes that invoked the programs that issued the secondary transactions.

If you had program-to-program switches during conversational transaction processing, these will be included in the lists.

| MESSAGES--PROGRAM TO PROGRAM |                | DATE 06/05/07 |  |
|------------------------------|----------------|---------------|--|
| DESTINATION                  | TOTAL MESSAGES |               |  |
| ELEANOR                      | 1              |               |  |
| SW1050                       | 1              |               |  |
| T2741N1                      | 1              |               |  |
| T2742N3                      | 1              |               |  |

| MESSAGES--PROGRAM TO PROGRAM |                | DATE 06/05/07 |  |
|------------------------------|----------------|---------------|--|
| TRANSACTION CODE             | TOTAL MESSAGES |               |  |
| TA10                         | 107            |               |  |

## Obtaining counts of unsent messages

The two reports titled Messages—Queued-But-Not-Sent summarize how many output messages were still in the message queues for interval covered by the input tapes.

The reports are illustrated in the following example. Command responses that were not sent to the terminal are indicated by (IMSSYS). The entry of NOTAVA indicates “no transaction available”. This would be the case if an output message were generated for an input not recorded in the system log input data or by a command input from the same terminal not recorded.

```

                MESSAGES - QUEUED BUT NOT SENT    DATE 06/05/07    PAGE 00001
                TOTAL
DESTINATION    MESSAGES
CTRLA01        1
SEGUNDO        37

```

## Auditing critical transactions

You can use the optional Messages report produced by the Statistical Analysis utility with the DFSIST40 program to examine the input and output data for specific transaction codes in detail. This allows you to audit exactly what was in an input message and possibly examine the output content for errors.

The report is illustrated in the following example.

```

MESSAGES
INPUT SEG=001 LEN=020* CONV12V0 Y c *
INPUT SEG=002 LEN=023*MESSAGE TO CONV12V0 > *
PGM SW SEG=001 LEN=066* CONV21V0 Y THIS DATA INSERTED TO THE SPA OF CONV21V0 c *
PGM SW SEG=002 LEN=042*MESSAGE TO CONV21V0 > *
OUTPUT SEG=001 LEN=066* Y THIS TERMINATED TO THE SPA OF CONV21V0 c *
OUTPUT SEG=002 LEN=035*RESPONSE TO INPUT TERMINAL > *
INPUT TRANSACTION SEQ OUTPUT SEQ
PREFIX CODE LTERM NO DATE TIME PREFIX LTERM NO DATE TIME
CONV12V0 TNDSE 00001 07.053 12.37.18 TNDSE 00001 07.053 12.37.19

```

## Log transaction analysis utility reports

You can obtain detailed data at the individual transaction level by using the Log Transaction Analysis utility (DFSILTA0). Although the data is not summarized by this utility, the detail report lines bring together many information items that help you assess the service given to a transaction type and the effect of the scheduling algorithm. The report shows actual response data because input data is the IMS log.

If you do not process the entire IMS log, data is presented from a starting checkpoint to a cutoff point. You limit the sample of transaction processing to be analyzed by specifying start time and duration in minutes, or you can give the number of checkpoints to be included after the starting checkpoint. Canceled messages are omitted.

The format of the Log-Analysis report and the full list of data items for each report detail line are shown in “Examining scheduling activity” on page 829. Times are given to the nearest tenth of a second and are elapsed times. You can see that the Processing Types field is a key description item.

Using the starting position and lengths of the fields in the report detail records you can specify a sort order for the second step in the utility execution.

The sort control statement to cause a report to be sequenced by message class and transaction priority is:

```
SORT FIELDS = (18,3,CH,A,16,1,CH,A)
```

You can also use the option of creating a DASD data set of the detail report records. Your installation could then develop an analysis program to extract and summarize the data.

## Examining scheduling activity

Using the data extracted by the Log Transaction Analysis utility you can examine the effect of your scheduling algorithm.

Each occurrence of a transaction primarily indicates:

- Message priority and message class
- Time in input queue
- Time to process
- Time in output queue
- Total time-in-system (measured between completion of message queue input to retrieval for output)

You can look at the processing type 'S' entries to see the send and receive times. You can sort the detail report lines by transaction code and look at any critical transactions requiring rapid response time. Refer to the following example for details on each table line item.

```
12.50.11 JOB00386 $HASP373 DFSILTA4 STARTED - INIT 4 - CLASS K - SYS STL1
12.50.12 JOB00386 SMF0001 DFSILTA4 ILTA DFSILTA0 0000
12.50.12 JOB00386 $HASP395 DFSILTA4 ENDED
----- JES2 JOB STATISTICS -----
05 JUN 2007 JOB EXECUTION DATE
30 CARDS READ
143 SYSOUT PRINT RECORDS
0 SYSOUT PUNCH RECORDS
10 SYSOUT SPOOL KBYTES
0.01 MINUTES EXECUTION TIME
1 //DFSILTA4 JOB 'TERRY',CLASS=K,MSGCLASS=A,MSGLEVEL=(1,1), JOB00386
// REGION=OM,TIME=1440,
// USER=USRT001,PASSWORD=
/*ROUTE PRINT THISCPU/IMSTST45
2 //JOBLIB DD DSN=IMSTESTL.TNUC0,DISP=SHR
3 // DD DSN=IMBBLD.I10RTS17.CRESLIB,DISP=SHR
4 //ILTA EXEC PGM=DFSILTA0
5 //HEADING DD SYSOUT=A
6 //PRINTER DD SYSOUT=A
7 //SYSUDUMP DD SYSOUT=A
8 //REPORT DD DUMMY
9 //TITLE DD *
//*
//* LOG INPUT FOLLOWS...
//*
10 //LOGIN0A DD DSN=IMSTESTL.SLDSP.IMSA.TYTY,DISP=SHR,
// UNIT=SYSDA,VOL=SER=DSHR03
11 //LOGIN01 DD DSN=IMSTESTL.SLDSP.IMS1.TYTY,DISP=SHR,
// UNIT=SYSDA,VOL=SER=DSHR03
12 //LOGIN02 DD DSN=IMSTESTL.SLDSP.IMS2.TYTY,DISP=SHR,
// UNIT=SYSDA,VOL=SER=DSHR03
13 //LOGIN03 DD DSN=IMSTESTL.SLDSP.IMS3.TYTY,DISP=SHR,
// UNIT=SYSDA,VOL=SER=DSHR03
ICH70001I USRT001 LAST ACCESS AT 12:50:09 ON TUESDAY, JUNE 5, 2007
IEF236I ALLOC. FOR DFSILTA4 ILTA
IEF237I 04A1 ALLOCATED TO JOBLIB
IEF237I 04B4 ALLOCATED TO
IEF237I JES2 ALLOCATED TO HEADING
IEF237I JES2 ALLOCATED TO PRINTER
IEF237I JES2 ALLOCATED TO SYSUDUMP
IEF237I DMY ALLOCATED TO REPORT
IEF237I JES2 ALLOCATED TO TITLE
IEF237I 04A9 ALLOCATED TO LOGIN0A
IEF237I 04A9 ALLOCATED TO LOGIN01
IEF237I 04A9 ALLOCATED TO LOGIN02
IEF237I 04A9 ALLOCATED TO LOGIN03
IEF237I 04A9 ALLOCATED TO LOGIN04
IEF237I 04A9 ALLOCATED TO LOGIN05
IEF237I 04A9 ALLOCATED TO LOGIN036
IEF142I DFSILTA4 ILTA - STEP WAS EXECUTED - COND CODE 0000
IEF285I USRT001.DFSILTA4.JOB00386.D0000102.? SYSOUT
IEF285I USRT001.DFSILTA4.JOB00386.D0000103.? SYSOUT
IEF285I USRT001.DFSILTA4.JOB00386.D0000104.? SYSOUT
IEF285I USRT001.DFSILTA4.JOB00386.D0000101.? SYSIN
IEF285I IMSTESTL.SLDSP.IMSA.TYTY KEPT
IEF285I VOL SER NOS= DSHR03.
IEF285I IMSTESTL.SLDSP.IMS1.TYTY KEPT
IEF285I VOL SER NOS= DSHR03.
IEF285I IMSTESTL.SLDSP.IMS2.TYTY KEPT
IEF285I VOL SER NOS= DSHR03.
```

```

IEF285I  IMSTESTL.SLDSP.IMS3.TYTY          KEPT
IEF285I  VOL SER NOS= DSHR03.
IEF285I  IMSTESTL.SLDSP2.IMS1.TYTY          KEPT
IEF285I  VOL SER NOS= DSHR03.
IEF285I  IMSTESTL.SLDSP2.IMS2.TYTY          KEPT
IEF285I  VOL SER NOS= DSHR03.
IEF285I  IMSTESTL.SLDSP2.IMS3.TYTY          KEPT
IEF285I  VOL SER NOS= DSHR03.
IEF373I  STEP/ILTA /START 2007156.1250
IEF374I  STEP/ILTA /STOP 2007156.1250 CPU    OMIN 00.02SEC SRB    OMIN 00.01SEC VIRT    16K SYS    296K EXT    808K SYS    11924K
IEF285I  IMSTESTL.TNUC0                      KEPT
IEF285I  VOL SER NOS= USER01.
IEF285I  IMSBLD.I10RTS17.CRESLIB            KEPT
IEF285I  VOL SER NOS= MVSS16.
IEF375I  JOB/DFSILTA4/START 2007156.1250
IEF376I  JOB/DFSILTA4/STOP 2007156.1250 CPU    OMIN 00.02SEC SRB    OMIN 00.01SEC
SPECIFIED START TIME IS 16:53:47.3
TO END OF FILE PROCESSED 9 TRANSACTIONS THIS RUN
REPORT CONTAINS 14 COMPLETE RECORD SETS
IMS LOG DATA FOR IMS3 STARTS AT TIME = 12:47:41.8, DATE = 2007.156
  JOB      STEP      DR      CLASSES
  NAME     NAME     ID *****
MPP1      MPP        1    1    2    3    4
MPP2      MPP        1    1    2    3    4
MPP3      MPP        1    1    2    3    4
IMS LOG DATA FOR IMS2 STARTS AT TIME = 12:46:39.5, DATE = 2007.156
IMS LOG DATA FOR IMS1 STARTS AT TIME = 12:34:26.8, DATE = 2007.156
IMS LOG DATA FOR IMSA STARTS AT TIME = 12:36:56.8, DATE = 2007.156
SEQ  TRANS  P  C  ***IN***  ****OUT**  P  PGM  DR  SMB*ENQ  MSG*SCHD  CNT*ENQ  MSG*END  CNT*GU  SYS  IN Q  PROC  OUT Q  TOTAL
NBR   CODE  R  L  LTERM   LTERM    T  NAME  ID  ID      HHMSST  HHMSST    HHMSST    HHMSST  ID  SSSST  SSSST  SSSST  SSSST
-----
00001  CHKPT  0001*****
00002  CHKPT  0001*****
00003  CHKPT  0001*****
00004  CHKPT  0001*****
00005          CTRLA02  CTRLA02  M
00006          CTRL    CTRL    M
00007          CTRLA02  CTRLA02  M
00008          CTRL    CTRL    M
00009          CTRL    CTRL    M
00010          CTRL    CTRL    M
00011  TRAN21V0 1  1  L62MVS1  L62MVS1  S  PGM2V0  1 1242411 1242411 1242413 1242413 1242414 131
00012  SMQR21C0 1  1  L62MVS1  L62MVS1  S  PGM2C0  1 1242419 1242419 1242421 1242421 1242424 131
00013          CTRLA01  CTRLA01  M
00014          CTRLA01  CTRLA01  M
00015  SMQR21C0 1  1  L62MVS1  L62MVS1  S  PGM2C0  1 1243147 1243147 1243149 1243149 1243150 131
00016  SMQR21C0 1  1  L62MVS1  L62MVS1  T  PGM2C0  1 1243279 1243279 1243281
00017          CTRL    CTRL    M
00018  CONV21V0 1  1  L62MVS1  L62MVS1  D  CPGM2V0  1 1243498 1243499 1243506
00019  CONV21V0 1  1  L62MVS1  L62MVS1  C  CPGM2V0  1 1243545 1243545 1243546 1243546 1243546 133
00020  CONV21V0 1  1  L62MVS1  L62MVS1  C  CPGM2V0  1 1243577 1243577 1243577 1243577 133
00021  CONV21V0 1  1  L62MVS1  L62MVS1  C  CPGM2V0  1 1244049 1244049 1244049 1244049 133
00022  CONV21V0 1  1  L62MVS1  L62MVS1  C  CPGM2V0  1 1244080 1244080 1244081 1244081 133
00026  CHKPT  0002*****
00027  CHKPT  0002*****
00028  CHKPT  0002*****
00029  CHKPT  0002*****
1234270 1234270 1234270 1234270 1234270 1
1236569 1236569 1236569 1236569 1236569 A
1239165 1239165 1239165 1239165 1239165 2
1240111 1240111 1240111 1240111 1240111 3
1241134 1241135 2 2
1241260 1241263 3 3
1241398 1241401 2 2
1241523 1241526 3 3
1242176 1242179 3 3
1242240 1242244 3 3
1242411 1242413 131
1242419 1242421 131
1242582 1242585 1 1
1243136 1243138 1 1
1243147 1243149 131
1243279 1243281 13
1243450 1243453 3 3
1243506 13
1243546 1243546 133
1243577 1243577 133
1244049 1244049 133
1244081 1244081 133
1248342 1248342 1248342 1248342 A
1249068 1249068 1249068 1249068 1
1249365 1249365 1249365 1249365 2
1250062 1250062 1250062 1250062 3

```

Table 85. Log Analysis report line format

| Identification                             | Starting position | Length | Note |
|--------------------------------------------|-------------------|--------|------|
| Sequence Number                            | 1                 | 5      | 1    |
| Transaction Code                           | 7                 | 8      |      |
| Priority of Transaction (PR)               | 16                | 1      |      |
| Class of Transaction (CL)                  | 18                | 3      |      |
| Input LTERM name                           | 22                | 8      |      |
| Output LTERM name                          | 33                | 8      |      |
| Processing Type (PT)                       | 44                | 1      | 2    |
| Program Name                               | 46                | 8      |      |
| Dependent Region ID                        | 55                | 3      |      |
| Time of SMB Enqueue (Transaction received) | 59                | 7      | 3    |
| Time of Message Schedule or GU             | 68                | 7      | 3    |

Table 85. Log Analysis report line format (continued)

| Identification                                     | Starting position | Length | Note |
|----------------------------------------------------|-------------------|--------|------|
| Time of CNT Enqueue (Message put on output queue)  | 77                | 7      | 3    |
| Time of Program End or Next Message GU             | 86                | 7      | 3    |
| Time of CNT GU (Output message starts to terminal) | 95                | 7      | 3    |
| System IDs (from LOGINxxx DD statement)            | 103               | 3      |      |
| Time in Input Queue                                | 106               | 6      | 4, 5 |
| Time Processing                                    | 113               | 6      | 5, 6 |
| Time in Output Queue                               | 120               | 6      | 5, 7 |
| Total Time                                         | 127               | 6      | 5, 8 |

**Key to the table:**

- Starting position 1 is a carriage control character which alters the starting positions of the fields when producing a report on disk.
- Processing types:
  - A** Abended transaction
  - C** Conversational Send/Receive Processing
  - D** Transmit Only Conversational Processing
  - F** /FORMAT entered (Transaction Code Field has MODNAME)
  - M** Message Switch
  - O** Region Occupancy (A region is occupied by a program that is processing transactions that existed in the input queue before the start checkpoint has encountered or a program scheduled by an unrecoverable message.)
  - P** Program Switch Send/Receive Processing
  - Q** Transmit Only Program Switch Processing
  - S** Send/Receive Processing
  - T** Transmit Only Processing
  - X** Conversational Program Switch, Send/Receive Processing
  - Y** Transmit Only Conversational Program Switch Processing
- Time HHMMSSST
- Input queue time is from SMB enqueue to message schedule.
- Time SSSST or OVRFLW (If the total seconds exceeds the field size, OVRFLW is printed).
- If the wait-for-input (WFI) system option is used, the time processing field also includes the wait time between transactions.
- Output queue time is from CNT enqueue to CNT GU.
- Total Time is from SMB enqueue to CNT dequeue. The total time spans the complete transaction.

---

## IMS accounting information

The nature of accounting methods varies a great deal among data processing installations. The IMS Transaction Manager presents special difficulties, because many and varied transactions are processed by a partnership of the control region and dependent regions.

Further, operationally discrete applications can be served concurrently.

For installations with IMS-dedicated processors, the overall cost of hardware and support functions is often charged back based on predicted and actual use by contributing groups. For shared systems, the processor usage can be the base for proportional cost.

Although IMS does not have an explicit accounting function, the individual events that make up the processing activity are recorded on the IMS log in considerable detail. Analysis of IMS log records can be used as a basis for charge-back algorithms. You can, for example, obtain a report from the Statistical Analysis utility of the number of transactions and the average number of DL/I calls for each transaction.

Another source of resource usage figures is the reports produced as a result of IMS Monitor data collection. Samples of processing activity can be taken on a regular basis. Accounting algorithms can use, for example, processor utilization by program.

Both of these approaches require further manipulation and calibration of the resource indicators.

If the DL/I address space option is used (LSO=S), accounting procedures based on SMF data will be affected. SMF statistics for IMS system data sets and Fast Path databases are accounted to the control region procedure. Full function databases are accounted to the DL/I address space procedure.

### Using the Application Accounting report

The Statistical Analysis utility produces an Application-Accounting report that you can use to assess machine charges. The following breakdown is provided for each transaction and for each program:

- The number of messages with related total and average processor time in seconds
- The number and type of DL/I message calls
- The number and type of DL/I database calls

The following example illustrates the output format.

| APPLICATION ACCOUNTING REPORT |             |          |        |    |    |      |    |    |      |     |     |        | DATE 06/05/07 |      |      |     | PAGE 00005 |      |        |      |
|-------------------------------|-------------|----------|--------|----|----|------|----|----|------|-----|-----|--------|---------------|------|------|-----|------------|------|--------|------|
| PROGRAM                       | TRANSACTION | MESSAGE- | COUNTS |    |    | DATA |    |    | BASE |     |     | COUNTS |               |      | CC   | OR  | RC         | TOT  | PROG   | AVG  |
| NAME                          | CODE        | PRI      | QTY    | GU | GN | ISRT | GU | GN | GNP  | GHU | GHN | GHNP   | ISRT          | DLET | REPL | NOT | 0          | CPU  | TIME   | TIME |
| CPGM2V0                       | CONV21V0    | 1        | 5      | 5  | 10 | 9    | 0  | 0  | 0    | 0   | 0   | 0      | 0             | 0    | 0    | 0   | 0          | 0.0S | 0.000S |      |
| PGM2C0                        | SMQR21C0    | 1        | 3      | 3  | 3  | 2    | 0  | 0  | 0    | 0   | 0   | 0      | 0             | 0    | 0    | 0   | 0          | 0.0S | 0.000S |      |
| PGM2V0                        | TRAN21V0    | 1        | 1      | 1  | 1  | 1    | 0  | 0  | 0    | 0   | 0   | 0      | 0             | 0    | 0    | 0   | 0          | 0.0S | 0.000S |      |

ICE143I

### Using IMS transaction profiles

You can use a composite picture of each transaction as a basis for estimating usage. The IMS Transaction profile can contain DL/I call requirements by type of call and

possibly items derived from path length for other processing blocks. You can weight the message count to allow for heavy DL/I use by the transaction. Transaction statistics can be obtained on a regular basis from /DISPLAY output, for example at end-of-day or before shutdown.

The profiles should characterize the IMS workload in such a way that growth trends and major deviations from the predicted load can be traced to the transaction codes responsible.

---

## Analyzing log records using KBLA utilities

Interpreting and analyzing IMS log data can be time consuming, complex, and error prone. To ensure that the resulting records are pertinent to a particular IMS image or that they correspond to a particular time period (without gaps) requires detailed setup and research.

IMS provides the Knowledge-Based Log Analysis (KBLA) utilities that simplify the interpretation and analysis of IMS log data. The KBLA utilities:

- Allow you to extract an interpreted version of IMS log records to produce easy-to-read output.
- Allow you to produce output that contains log records which are grouped by an attribute or a resource, such as a unit of work.
- Simplify JCL stream creation to help prevent JCL errors. The necessary control statements are created in a syntax-correct form, so that you can focus on diagnosis rather than on how to extract log records.
- Provide a simple, quick, and automated method to:
  - Sort a set of log data sets
  - Verify the completeness of a list of logs (created manually or by DBRC)
  - Sort the logs into the proper sequence by IMS image
  - Determine where log data sets might be missing in the sequence
  - Determine the number of missing records and the time range that is missing to help identify the missing logs
- Reduce problem turnaround time with IBM Support by reducing the number and the size of log data sets, and by ensuring the appropriateness of log data sets that are sent to Support.

### Related concepts:

“IMS Application Menu” on page 119

“Using IMS system log utilities” on page 349

## Overview of KBLA routines

The KBLA log-formatting routines help you examine and display data from IMS log data sets.

### Knowledge-Based Formatting Print routine (DFSKBLAK)

Provides a clear, simple description of each event represented by a log record, including the meaning of the various fields and flags.

### Knowledge-Based Summary Formatting Print routine (DFSKBLAS)

Prints and displays the header and description of a log record, but not the record itself. This information is useful when you need a general understanding of the log records associated with a particular resource.

### Knowledge-Based Basic Formatting Print routine (DFSKBLA3)

This routine is similar in function to the Record Format and Print Module



(DFSERA30), but its output also provides a brief description of the log record identifier (or log record type). This information can be useful if you are unfamiliar with the record types present in an IMS log. The Knowledge-Based Basic Formatting Print routine also interprets the prefix fields for the IMS type X'01' and type X'03' log records.

The following KBLA routines help with log analysis and diagnosis:

#### **DBCTL Transaction Analysis utility (DFSKDBC0)**

The DBCTL Transaction Analysis utility (DFSKDBC0) performs the functions of both the Log Transaction Analysis utility (DFSILTA0) and the Fast Path Log Analysis utility (DBFULTA0). The DBCTL Transaction Analysis utility also performs a sort of the data using key fields that can help to identify potential performance problems. The DBCTL Transaction Analysis utility uses the X'07' and X'5937' log records or the X'5938' log record to gather statistics.

#### **Deadlock Trace Record Analysis utility (DFSKTDL0)**

The Deadlock Trace Record Analysis utility is used to format and summarize data extracted from X'67FF' pseudo-abend records for database-related deadlocks. These records are identified by a four-character requester identification that indicates a pseudo-abend condition, and an eight-character element identification that indicates deadlock.

The Deadlock Trace Record Analysis utility processes the logical deadlock events and the individual deadlocks to produce reports that:

- Provide details about the deadlock
- Show the hierarchy of the participants in the deadlocks (relative to the victim in each deadlock)
- Summarize the deadlock activity, including deadlocks by:
  - Hour
  - DBMS
  - State
  - Lock type
  - DBD
  - PSB
  - Lock name
  - RBA

#### **IMS Records User Data Scrub utility (DFSKSCR0)**

The IMS Records User Data Scrub utility (DFSKSCR0) performs a scan of all of the IMS logs. The utility deletes those record parts that might contain sensitive or confidential customer business transaction information, in particular, X'01', X'03', X'4002', X'4082', X'4086', X'50', X'5901', X'5903', X'5920', X'5950', and X'67'. The deletion does not compromise the integrity or the content of the vital IMS system data. This scrub is useful when you must send IMS log data to an outside organization for analysis.

#### **IRLM Lock Trace Analysis Utilities (DFSKLTA0, DFSKLTB0, DFSKLTC0)**

IRLM Lock Trace Analysis consists of three programs (DFSKLTA0, DFSKLTB0, and DFSKLTC0) that run serially to perform IRLM Lock Trace Analysis. These utilities create and combine several outputs, including one based on wait time order. In the wait-time-ordered report, databases are ordered by the total lock wait time during the trace.

DFSKLTA0 is run first to create the control file of global data management block (DMB) numbers and their respective database names. DFSKLTB0 is

then used to create the lock wait detail and summary records. DFSKLTC0 formats and prints the information and creates the optional output data set.

The IRLM Lock Trace Analysis utilities can be useful in finding database or application issues that impact transaction response times by causing frequent and long lock waits.

You can save the report to a data set, which you can then sort using either the knowledge-based log analysis ISPF tools or an editor's SORT command. The IRLM Lock Trace Analysis panel provides some predefined sort options, such as Sorted by Database Name and Sorted by RBA.

#### **Log Record Processing Rate Analysis utility (DFSKRSR0)**

The Log Record Processing Rate Analysis utility (DFSKRSR0) is used to generate reports that summarize the volume of log data that is being generated by an IMS subsystem. The volume of log data is expressed in number of records per second and number of bytes per second. The detailed log record processing rate data is broken down by log record type, and by subtype within record type, if requested.

The Log Record Processing Rate Analysis utility can be used to:

- Determine the size of log data sets or archiving frequency
- Track data volume for purposes such as Remote Site Recovery (RSR) planning

Selection criteria for the DFSKRSR0 trace table entries can include:

- Log record type
- Subtypes within log record types
- Time range data
- Log sequence number ranges
- Number of records to be processed or skipped

#### **Log Summary utility (DFSКСUM0)**

The Log Summary utility (DFSКСUM0) can create a summary of the content of the log data sets, characteristics of the IMS system that produced the log data (when statistical log records are present), and some statistical information related to transactions, programs, and databases. The Log Summary utility also includes the following information:

- Input IMS logs used for the utility execution
- IMSID
- First and last Log Sequence Number (LSN) in the log
- Time stamp (UTC) and local time of the first and last log record
- Difference between UTC and local time (HHMM format)
- Elapsed time on selected logs
- Total number of log records in the log data set
- Presence of internal trace record, system restarts, dump log record, and system checkpoint
- Number of log records present for each record ID
- Statistics related to individual transactions, programs, databases, and other specifiable resources
- Gaps in the log

You can use this utility to specify search criteria to extract the supported log records. Log record types that can be processed during the analysis

include: X'01', X'03', X'07', X'08', X'11', X'12', X'13', X'16', X'20', X'21', X'27', X'29', X'31', X'33', X'35', X'36', X'36', X'37', X'38', X'3F', X'40', X'41', X'42', X'42', X'45', X'47', X'4C', X'50', X'51', X'52', X'56', X'57', X'59', X'63', X'70', and X'72'.

Search criteria can include the following:

- UOW (unit of work)
- LTERM name
- NODE name
- TRANSACTION name
- Program name (PGM)
- DBD name
- AREA name
- User ID
- RBA
- PST
- Recovery Token
- Message device relative record number (DRRN)
- A generic character or hexadecimal search string

The “Dynamic Search Keys” parameter allows you to create an increasingly refined and broader search argument to include additional records that are related to the specified search criteria.

#### **MSC Link Performance Formatting utility (DFSKMSC0)**

The MSC Link Performance Formatting utility (DFSKMSC0) uses the IMS MSC link trace log records to provide information about link response times, which can be used to help isolate performance problems with MSC links. You can specify either of the following types of analysis:

**Detail** The Detail Selection routine (DFSKBST0) produces a report that contains the individual response times in milliseconds of every send data and receive data for each MSC link that is traced.

#### **Summary**

The Summary Selection routine (DFSKMSCD) produces a report that contains the average response time in milliseconds of the total number of send data and receive data values for each link trace.

The enhanced log formatting routines produce formatted output only if the IMS input logs contain X'6701' records generated using the /TRACE SET ON LINK*link#* command.

#### **RECON Query of Log Data Set Names utility (DFSKARC0)**

The RECON Query of Log Data Set Names utility (DFSKARC0) analyzes the RECON data sets to find appropriate log data set names. Based on the control statements that you provide, it can determine the following information:

- Data set names for OLDS, SLDS, or PRILOG data sets
- The names for the primary or secondary log data sets
- The volume serial numbers for these data sets

The logs can be selected by any or all of the following parameters:

- Starting date and time
- Ending date and time

- SYSID

Information from the RECON data sets can be extracted from either of the following places:

- The RECON data sets themselves, when available on the system that executed DFSKARCO
- A pre-existing report generated by the DBRC LIST.LOG command

#### **Statistic Log Record Analysis utility (DFSKDVS0)**

The Statistic Log Record Analysis utility (DFSKDVS0) processes the X'45' log records that are generated at each IMS checkpoint and formats a report that is a verbal description and interpretation for all of the fields contained in the X'45' log record, including its subcodes. This information can be used to look for bottlenecks within the IMS system or to detect trends in internal resource usage that can help to determine if tuning is necessary. You must use the /CHECKPOINT STATISTICS command to ensure that the X'45' log record contains the information required for this enhanced formatting.

#### **Trace Record Extract utility (DFSKXTR0)**

The need for information for problem diagnostics or accumulation of statistics drives trace table logging, which often generates a large number of records. The Trace Record Extract utility (DFSKXTR0) reads an IMS log data set (OLDS or SLDS) or trace data set (DFSTRAXX) to produce a subset of information that meets specific selection criteria for subsequent formatting by File Select and Formatting Print utility (DFSERA10) and DFSERA10-related exit routines. DFSKXTR0 reduces the amount of data unrelated to the selection criteria in the formatted reports.

Unlike DFSERA10, which searches entire log records for specific strings, DFSKXTR0 searches individual trace table entries using selection criteria. Trace table entries are relatively small, 32- or 64-byte entries. Although the data in the entries can vary depending on what is specifically being traced, the structure is a constant eight words of data.

Search criteria is entered as a paired specification identifying the word, half word, or byte to be examined, and the associated data value. Optionally, the entire log record containing the matching trace table entries can be extracted as well.

After searching the trace table entries, DFSKTXR0 stores the entries that match selection criteria into new log records. These new log records can subsequently be passed to the DFSERA60 exit routine for formatting.

Selection criteria for the trace table entries can include:

- Trace table IDs
- Character string searches at the trace entry level
- Specific words, half words or bytes in the trace entries
- Time range data
- Entire log records to be matched
- Number of records during processing
- Mode of SCAN to read the log records without actual processing

#### **PSB Database Update Analysis (DFSKPSB0)**

This option generates produces a statistical summary of the database updates that are performed by the PSB. Database update log records X'50' and X'59' are counted and summed for each PSB occurrence that has issued

database update calls. For full-function databases, counts are also broken down by call types: ISRT, DLET, REPL, ROLX. No comparable counts are made for Fast Path databases.

You can filter the log records that are extracted for processing by database name or PSB name, or by all database records that are processed.

You can sort the output data by database name, PSB name, or the number of analyzed records.

**Related reading:** For complete information about the individual KBLA utilities, see *IMS Version 11 System Utilities*.

## ISPF interface for KBLA

Knowledge-Based Log Analysis (KBLA) is a collection of IMS utilities that enable you to select, format, and analyze log records. Use the ISPF interface to create and run the jobs for various log-related utilities, and to access other ISPF applications.

KBLA uses an ISPF panel-driven user interface to simplify JCL job creation and to prevent JCL errors. KBLA generates the JCL and control statements that are necessary to run the supported utilities. This JCL preparation allows you to focus on the output of the utility used rather than on how to code JCL to extract information.

KBLA provides a set of ISPF panels that you can use to access log analysis-related utilities and functions. A set of help panels for each function or sub-function can also be displayed. You can invoke field level help by placing the cursor on the subject field and pressing the HELP Key (PF1). The following figure shows the structure of the KBLA ISPF panels.




Figure 95. KBLA ISPF panel structure

## Invoking KBLA

KBLA can be invoked in two ways. It can be invoked from commands issued on ISPF Option 6, from the IMS Application Menu, or it can be invoked directly.

To access KBLA from the IMS Application Menu:

1. Access the IMS Application Menu by typing the following command: EXEC 'hlq.SDFSEEXEC(DFSAPPL)' 'HLQ(hlq)'

**Note:** *hlq* is the high-level-qualifier that is associated with the installed IMS subsystem.

The IMS Application Menu panel is displayed.

```
IMS Application Menu

Select the desired application and press Enter.

1 Single Point of Control (SPOC)
2 Manage resources
3 Knowledge-Based Log Analysis (KBLA)
4 HALDB Partition Definition Utility (PDU)
5 Syntax Checker for IMS parameters (SC)
6 Installation Verification Program (IVP)
7 IVP Export Utility (IVPEX)
8 IPCS with IMS Dump Formatter (IPCS)
9 Abend Search and Notification (ASN)

To exit the application, press F3.

Command ==>
```

Figure 96. IMS Application Menu

2. From the IMS Application Menu, select option 3 for KBLA. The main KBLA menu opens:

```
Knowledge Based Log Analysis   IMS Version 11

Command ==>

TIME...17:15:07
DATE...2008/10/01
JULIAN..2008.274
USERID..TSOUSER

Select any of the following tasks and press ENTER .

Tasks . .   0. KBLA Environment Maintenance
              1. IMS Log Utilities
              2. IMS Log Formatting
              3. IMS Log Data Set Summary
              4. IMS Knowledge Based Analysis
              5. Log Selection
              6. User Supplied Utilities

To Exit the KBLA MAIN menu, press END .
For Help place cursor on any field and press PF1 .
```

Figure 97. Main panel for KBLA

Access the Main KBLA panel directly by typing the following command: EXEC 'hlq.SDFSEEXEC(DFSKBST)' 'HLQ(hlq)'

### Maintaining the KBLA environment with option 0

Use option 0, suboption 1 to perform global maintenance on your KBLA environment. Parameters entered in KBLA panels are stored as ISPF variables and are used by other panels, as appropriate. It can be useful to complete the KBLA Maintenance Environment panel first in order to tailor KBLA to specific needs.

Option 0 can also be used to delete work data sets that have been created by KBLA, or other data sets whose high-level qualifier matches a value you have specified on the panel.

If entries are not specified, KBLA fills some fields with default values. The default JCL job statement is extracted from hlq.SDFSRLIB(DFSKJ0BK). Other default values include:

*Table 86. KBLA fields and default values*

| Field                            | Default value |
|----------------------------------|---------------|
| KBLA Test Loadlib                | none          |
| IMS Version                      | 11            |
| IMS.SDFSRESL DSN                 | hlq.SDFSRESL  |
| Verify LOG DSN Exists            | Y             |
| Default Output Space Parms: Type | CYL           |
| Primary                          | 100           |
| Secondary                        | 50            |
| Number of vols                   | 1             |
| Route output to alternate system | N             |

Suboptions available for KBLA environment maintenance include:

1. Tailoring the KBLA panels according to your ISPF environment, DSN naming convention, JOB JCL requirements, IMS.SDFSRESL used, and RECON-related information
2. Creating a list of the work data sets that have been created as you use KBLA
3. Viewing the list of the work data sets
4. Deleting work data sets that are no longer needed

## Defining the selection of IMS logs using option 5

Use option 5 to create a list of logs to be processed by KBLA. KBLA can process logs from different sources. Option 5 can also be used to sort records within log data sets.

Sometimes multiple IMS log data sets must be used. The RECON data sets can be accessed to extract the names of PRIOLD, PRISLD, PRILOG, SECOLD, SECSLD, or SECLOG data sets. A copied or manually entered list of logs can also be used as input to KBLA. In addition to accepting a list of IMS logs, KBLA can separate the list by IMS SYSID into multiple lists, each sorted by time stamps.

The following suboptions create and sort the final list of logs for input to KBLA:

- Suboption 1: Extract log data sets from the RECON data sets or as the result of a LIST.LOG RECON command.
- Suboption 2: Create a member or members in your SDFSJKJCL PDS that contains a list of logs to be processed by KBLA options. The list of logs can be entered manually or pasted from another source.
- Suboption 3: Use the list of logs created in suboption 2 to create new sorted members in the SDFSJKJCL PDS, which can be used as input by other KBLA processing options. This suboption separates the log data set names by IMS SYSID and sorts the log data set names by time stamp or log sequence number.



- Suboption 4: Sort records within logs by using up to four criteria, or the log sequence number (LSN), or the time stamp contained in the log suffix. This suboption is used to process unformatted log records that have been previously selected by KBLA (use Option 2 with Log Formatting Type 'U'). After this suboption sorts individual log records within a log data set and sends the sorted log to KBLA for processing, the sorted log records can be used as input for the formatting and analysis steps of KBLA.

## Using KBLA to run a job against IMS log records

Use the remaining four options from the main panel of the Knowledge-Based Log Analysis utility to format and analyze a log record.

A brief description of each option follows.

### Option 1: IMS Log Utilities

This option allows you to access the following log analysis utilities:

1. Log Transaction Analysis utility (DFSILTA0)
2. Fast Path Log Analysis utility (DBFULTA0)
3. Statistical Analysis utility (DFSISTS0)
4. Log Merge utility (DFSMTMG0)
5. Log Recovery utility (DFSULTR0)
6. Program isolation trace record format and print module (DFSERA40)
7. IMS Records User Data Scrub utility (DFSKSCR0)

### Option 2: IMS Log Formatting

This option allows you to extract or format log records. Suboptions include:

1. IMS Resources Formatting
2. IMS Subcomponent Log Filtering
3. KBLA Log Records Formatting
4. IMS Trace Formatting
5. Snap/Pseudo-Abend Record Formatting
  - a. Deadlock Trace Record Analysis utility (DFSKTDL0) for K-formatted database deadlock traces
  - b. DFSERA10 using the Record format and print module (DFSERA30) for all other traces
6. Free-Form Log Filtering

A range selection panel can be invoked from KBLA Options 2.1, 2.2, 2.3, and 2.6 from which log records can be filtered by the following ranges:

- Time stamp
- Log sequence number (LSN)
- Recovery tokens
- /LOG command strings

A number of records to skip and process can also be specified.

### Option 3: IMS Log Data Set Summary

This option produces a summary of the log data set that you specified along with some statistical information. The summary function includes:

- First and last line sequence number (LSN) in the log
- Time stamp (UTC) of the first and last log record
- Total number of log records in the log data set
- Presence of internal trace records, system restarts, dump log records, system checkpoint
- Number of log records present for each record ID
- System configuration
- Transaction, program and data bases records instances

### Option 4: IMS Knowledge-Based Analysis

This option provides different methods for knowledge-based formatting, analysis and interpretation of specific IMS log records. You can access different utilities with the following suboptions, each one with its own ISPF panel interface.

1. Log Summary utility (DFSKSUM0)  
Use DFSKSUM0 to create summaries of log records, detailed reports based on specific search criteria, and copies of log records that have been processed.
2. MSC Link Performance Formatting utility (DFSKMSC0)  
Use DFSKMSC0 to generate reports about link response times.
3. Statistic Log Record Analysis utility (DFSKDVS0)  
Use DFSKDVS0 to process the X'45' log records that are generated at each IMS checkpoint and to format a report showing the detailed statistic information between each pair of checkpoints.
4. Trace Record Extract utility (DFSKXTR0)  
Use DFSKXTR0 to read an IMS log data set (OLDS/SLDS) or trace data set (DFSTRAx) to produce a subset of information that meets specific selection criteria from among the large number of records generated during problem diagnostics or accumulation of statistics.
5. IRLM Lock Trace Analysis utilities (DFSKLTA0, DFSKLTB0, DFSKLTC0)  
Use DFSKLTA0, DFSKLTB0, and DFSKLTC0 to identify database issues or application issues that impact transaction response times by causing frequent and long lock waits.
6. DBCTL Transaction Analysis utility (DFSKDBC0)  
Use DFSKDBC0 to combine some of the information in the DBFULTA0 and DFSILTA0 utilities with information that is specific to DBCTL.
7. Log Record Processing Rate Analysis utility (DFSKRSR0)  
Use DFSKRSR0 to generate reports that summarize the volume of the log data that is being generated by an IMS subsystem.
8. Database Error Pointer Analysis  
This panel provides an effective way to analyze large quantities of log and trace data related to common database pointer error problems. Application abends, utility, and pointer checker reports can provide diagnostic information identifying the DBD, DSID, BLKSZ, data set type (VSAM or OSAM), error pointers, and date/time period related to the pointer errors. Information entered on this panel is used to generate JCL job streams to extract records from log and trace data sets. This selection combines the use of the File Select

and Formatting utility (DFSERA10) to extract X'50' and X'59' database update log records, and the Trace Record Extract utility (DFSXTR0) to extract X'67FA' DL/I and lock trace record entries.

#### 9. PSB Database Update Analysis

This option generates produces a statistical summary of the database updates that are performed by the PSB. Database update log records X'50' and X'59' are counted and summed for each PSB occurrence that has issued database update calls. For full-function databases, counts are also broken down by call types: ISRT, DLET, REPL, ROLX. No comparable counts are made for Fast Path databases.

You can filter the log records that are extracted for processing by database name or PSB name, or by all database records that are processed.

You can sort the output data by database name, PSB name, or the number of analyzed records.

The 16-character recovery token determines the affinity of the records. The records can be grouped by the first 12 characters of the recovery token, which represents a PSB in the output by a single record, or by all 16 characters, which represents each checkpoint individually in the output.

You can further filter the log records by specified time ranges or by the number of records to skip before processing and the number of records to process can be specified.

Processing consists of the following steps:

- a. Optional extraction of database update records from the log by specified time ranges or skip or process values
- b. Optional extraction of database update records by specified PSB names or database names
- c. Sorting and grouping of database update records by the following criteria:
  - PSB name
  - Recovery token
  - PST number
  - Database name
- d. Counting of the database records by recovery token affinity
- e. Optional sorting of the output by specified criteria
- f. Optional creation of a data set that contains the sorted filtered log records

### **User-Supplied utilities by using option 6**

Use option 6 to access any external log processing utility or IMS tool available in your environment.

You can launch up to 4 different utilities or tools from this panel by specifying the corresponding EXEC statement in the TSO Exec field.

---

## Part 6. Appendixes



---

## Notices

This information was developed for products and services offered in the U.S.A.

This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive, MD-NC119  
Armonk, NY 10504-1785  
United States of America*

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan, Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
J46A/G4  
555 Bailey Avenue  
San Jose, CA 95141-1003  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to

IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS," without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. \_enter the year or years\_. All rights reserved.

## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com/legal/copytrade.shtml)<sup>®</sup> are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

---

## Programming interface information

This information documents Product-sensitive Programming Interface and Associated Guidance Information and General-use Programming Interface and Associated Guidance Information.

Product-sensitive Programming Interfaces allow the customer installation to perform tasks such as diagnosing, modifying, monitoring, repairing, tailoring, or tuning of this software product. Use of such interfaces creates dependencies on the detailed design or implementation of the IBM software product. Product-sensitive Programming Interfaces should be used only for these specialized purposes. Because of their dependencies on detailed design and implementation, it is to be expected that programs written to such interfaces may need to be changed in order to run with new product releases or versions, or as a result of service.



Product-sensitive Programming Interface and Associated Guidance Information is identified where it occurs, either by an introductory statement to a section or topic, or by a Product-sensitive programming interface label. IBM requires that the preceding statement, and any statement in this information that refers to the preceding statement, be included in any whole or partial copy made of the information described by such a statement.

General-use programming interfaces allow the customer to write programs that obtain the services of IMS. General-use Programming Interface and Associated Guidance Information is identified where it occurs, either by an introductory statement to a section or topic or by a General-use programming interface label.

---

## Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

---

## Bibliography

This bibliography lists all of the publications in the IMS Version 11 library, supplemental publications, publication collections, and accessibility titles cited in the IMS Version 11 library.

### IMS Version 11 library

| Title                                                                     | Acronym | Order number |
|---------------------------------------------------------------------------|---------|--------------|
| <i>IMS Version 11 Application Programming</i>                             | APG     | SC19-2428    |
| <i>IMS Version 11 Application Programming APIs</i>                        | APR     | SC19-2429    |
| <i>IMS Version 11 Commands, Volume 1: IMS Commands A-M</i>                | CR1     | SC19-2430    |
| <i>IMS Version 11 Commands, Volume 2: IMS Commands N-V</i>                | CR2     | SC19-2431    |
| <i>IMS Version 11 Commands, Volume 3: IMS Component and z/OS Commands</i> | CR3     | SC19-2432    |
| <i>IMS Version 11 Communications and Connections</i>                      | CCG     | SC19-2433    |
| <i>IMS Version 11 Database Administration</i>                             | DAG     | SC19-2434    |
| <i>IMS Version 11 Database Utilities</i>                                  | DUR     | SC19-2435    |
| <i>IMS Version 11 Diagnosis</i>                                           | DGR     | GC19-2436    |
| <i>IMS Version 11 Exit Routines</i>                                       | ERR     | SC19-2437    |
| <i>IMS Version 11 Installation</i>                                        | INS     | GC19-2438    |
| <i>IMS Version 11 Licensed Program Specifications</i>                     | LPS     | GC19-2439    |
| <i>IMS Version 11 Master Index and Glossary</i>                           | MIG     | SC19-2440    |
| <i>IMS Version 11 Messages and Codes, Volume 1: DFS Messages</i>          | MC1     | GC19-4232    |
| <i>IMS Version 11 Messages and Codes, Volume 2: Non-DFS Messages</i>      | MC2     | GC19-4233    |
| <i>IMS Version 11 Messages and Codes, Volume 3: IMS Abend Codes</i>       | MC3     | GC19-4234    |
| <i>IMS Version 11 Messages and Codes, Volume 4: IMS Component Codes</i>   | MC4     | GC19-4235    |
| <i>IMS Version 11 Operations and Automation</i>                           | OAG     | SC19-2441    |
| <i>IMS Version 11 Release Planning</i>                                    | RPG     | GC19-2442    |
| <i>IMS Version 11 System Administration</i>                               | SAG     | SC19-2443    |
| <i>IMS Version 11 System Definition</i>                                   | SDG     | GC19-2444    |
| <i>IMS Version 11 System Programming APIs</i>                             | SPR     | SC19-2445    |
| <i>IMS Version 11 System Utilities</i>                                    | SUR     | SC19-2446    |

### Supplementary publications

| Title                                                                                             | Order number |
|---------------------------------------------------------------------------------------------------|--------------|
| <i>IMS and SOA Executive Overview</i>                                                             | GC19-2516    |
| <i>Program Directory for Information Management System Transaction and Database Servers V11.0</i> | GI10-8788    |
| <i>IMS Version 11 Fact Sheet</i>                                                                  | GC19-2451    |
| <i>IRLM Messages and Codes</i>                                                                    | GC19-2666    |

## Publication collections

| Title                      | Format | Order number |
|----------------------------|--------|--------------|
| IMS Version 11 Product Kit | CD     | SK5T-7375    |

## Accessibility titles cited in the IMS Version 11 library

| Title                                  | Order number |
|----------------------------------------|--------------|
| <i>z/OS TSO/E Primer</i>               | SA22-7787    |
| <i>z/OS TSO/E User's Guide</i>         | SA22-7794    |
| <i>z/OS ISPF User's Guide Volume 1</i> | SC34-4822    |

---

# Index

## Special characters

//DFSSTAT report  
DD statements 811  
OSAM-Buffer-Pool Report 813  
PST-Accounting Report 811  
Sequential Buffering Detail Report 816  
Sequential-Buffering-Summary Report 814  
specifying in JCL 811  
types of reports 811  
VSAM-Buffer-Pool report 812  
/CHANGE SURVEILLANCE command  
XRF usage 597  
/CHECKPOINT command 75  
/DBDUMP command  
database backup copies 517  
/DISPLAY command  
for monitoring 375  
MODIFY 441, 445  
use in accounting 832  
/ERESTART BACKUP command 608, 621  
/MODIFY command  
PREPARE 450  
using for online change 441  
/NRESTART  
restart after IMS failure 512  
/OPNDST command 139  
/RMGENJCL command  
generating JCL 493  
/RMxxxxxx commands  
overview 482  
/SIGN ON/OFF Security exit  
routine 331  
/START AVM command 608  
/START command  
defining sharing level with 499  
/START IMS command 608  
/START SUBSYSTEM SSM command 13  
/START SURVEILLANCE command  
XRF usage 597  
/STOP SURVEILLANCE command  
XRF usage 597  
/TRACE command  
for program isolation 381  
IMS Monitor 378  
to turn on IMS Monitor 388  
/TRACE SET OFF command 32  
/TRACE SET ON command 32

## Numerics

37x5 Communication Controllers  
gateway considerations 643  
ISC link between active and alternate  
IMS system 597  
requirement for XRF 595

## A

abend codes  
Base Primitive Environment (BPE) 33  
abend formatting module  
binding 134  
DFSAMFD0 134  
uninstalling 135  
abnormal termination 76  
ACB (application control block)  
changing online 454  
library  
allocating 399  
changing members online 454  
member online change 454  
library data sets  
allocating dynamically 31  
dynamic allocation 31  
online change 454  
passwords for XRF 647  
placement for performance 412  
placing in 64-bit storage 31  
ACB library  
global online change 454  
inactive library  
resizing online 398  
member online change 50, 454  
ACB library member online change  
overview 454  
ACB member  
online change  
associating 448  
associating with DRD  
commands 448  
ACBLIB  
global online change 50  
member online change 50  
access control methods  
considerations 338  
access method services (AMS) 560  
REPRO command 556  
accessibility  
features xiv  
keyboard shortcuts xiv  
active IMS  
bringing up 608  
cycle of processing 624  
definition 665  
initializing entry in USERVAR  
table 601  
procedures at takeover 600, 613  
processing during tracking  
phase 588, 611  
processing of 600  
surveillance signals, sending 597, 611  
tracking by alternate IMS  
system 597, 611  
active IMS system  
definition 586  
processing of 600  
address space  
dispatching priority 372  
address space (*continued*)  
relative importance 372  
workload management 372  
address spaces  
IMS, in a CSL 39  
administration  
illustrated overview 3  
ODBA application programs  
security considerations 344  
system  
concepts 31  
system, overview 3  
advanced program-to-program  
communication/IMS 69  
affinity  
management, bypassing 297  
termination 293  
allocating  
OLDS  
defining OLDS 94  
SLDS 96  
WADS 95  
allocation  
database record 554  
log record 554  
logging in the RECON data set 497  
RECON data set 535  
allocation of data sets  
ACBLIB 399  
altering structures 179  
alternate IMS system  
active IMS, becoming 613  
allocating databases 597  
allocating IMS system log 610  
availability for non-XRF work 588  
backup sessions for class-1  
terminals 591  
closing backup sessions for class-1  
terminals 597  
considering a takeover 613  
cycle of processing 624  
databases, allocating 597  
definition 586, 665  
detecting problems in active  
IMS 597, 611  
executing new transactions at  
takeover 597, 617  
loading MSDBs 610  
logging on to 597, 601  
non-XRF workload  
at takeover 632  
available for 588  
opening backup sessions for class-1  
terminals 597  
opening databases 597  
performing I/O toleration 615  
planning workload on 588  
processing during a takeover 613  
processing during post-takeover  
phase 621

- alternate IMS system (*continued*)
  - processing during tracking phase 588, 611
  - starting up 608
  - synchronization with active IMS 610
  - system log, allocating 610
  - takeover, considering 613
  - termination, at 623
  - third system as 621
  - tracking the active IMS 597, 611
  - updating its control blocks 597, 611
- ALTRESL parameter 120
- AMS 556, 560
- AO (automated operator) application
  - program
    - description 74
- AO (automated operator) application programs
  - dynamic monitoring 375
  - ICMD call 315
  - security 315
- AO application program security
  - AOI= execution parameter 316
- AO applications
  - user ID substitutions 318
- AO command security
  - command authorization exit routine 320
  - implementing with RACF 319
  - initialization EXEC parameters 317
- AOI (Automated Operator Interface)
  - description 113
- AOI transactions
  - scheduling 230
- AOI= execution parameter
  - AO application program security 316
- APAR (Authorized Program Analysis Report)
  - using SMP/E 191
- APF authorization
  - for z/OS interface 125
  - required by z/OS 136
  - required for IRLM 137
- APPC (advanced program-to-program communication)
  - descriptor
    - as a TM resource 300
    - generic resource name, specifying 291
- APPC (advanced program-to-program communication)/IMS
  - description 69
  - security 326
  - shared-queues environment 228
    - asynchronous messages 228
    - synchronous messages 228
  - transaction profile names 85
- APPC/MVS administration dialog updates 137
- APPLCTN macro statement
  - used for online change 403
- Application Accounting report
  - example 832
  - use for accounting 832
- application change control, checklist 395
- application changes
  - introducing in an active IMS system 398
  - planning how they affect the system 395
- application control block (ACB)
  - changing online 454
- library
  - allocating 399
  - changing members online 454
  - member online change 454
- library data sets
  - allocating dynamically 31
  - dynamic allocation 31
- online change 454
- placement for performance 412
- placing in 64-bit storage 31
- application design
  - performance considerations 425
- application documentation 81
- application preload
  - guidelines 422
- application processing intent 237
- application program
  - commit point 92
  - data sets that support 31
  - design
    - performance considerations 425
  - design reviews 82
  - message-driven 72
  - sync point 92
- application program deadlock
  - with block-level sharing 241
- application programs
  - CPI-C driven
    - security considerations 321
  - INIT call 71
  - ODBA
    - security considerations 322
  - online changes 402
  - participating in data sharing 237
  - scheduling
    - against unavailable data 71
  - sensitivity to unavailable data 71
  - unavailable data
    - sensitivity to unavailable data 71
  - with XRF 595
- applications
  - changing
    - in an active IMS system 398
  - serial
    - in a shared queues environment 210
- APPLID name
  - in a generic resources environment 292
- APPLID= keyword
  - defining VTAM application name for XRF 647
- APSB SAF security
  - disabling 335
  - enabling 335
- archive, OLDS 74
- archiving
  - customizing 101
  - log records
    - automatic 99
- archiving (*continued*)
  - log records (*continued*)
    - customizing 101
    - manual process 99
  - OLDS 94
  - online log data sets
    - using DBRC 487
- area data set
  - initiating in the RECON data set 507
- area data sets
  - XRF 656
- area-level data sharing 233
- ARM 141
- assigning a sharing level with DBRC 499
- attention notice
  - installing preventive service
    - ACCEPT before APPLY 200
    - ACCEPT without APPLY 195
  - interface considerations 125
  - SYS1.NUCLEUS 131
  - VTAM interface considerations 139
- audit log
  - OM
    - record format 155
- audit trail 153, 154
- auditing operations 337
- auditing transactions 828
- authorization
  - changing database 572
  - database, changing 572
- authorizing CQS
  - connections 178
  - registration 178
- autologon terminal
  - shared queues 209
- automated operations
  - advantages 88
  - introduction 113
  - overview 88
  - what to automate 88
- automated operator applications
  - user ID substitutions 318
- Automated Operator Interface (AOI)
  - description 113
- automated procedures
  - single point of control 45
  - supported consoles 45
- automatic archiving 99
- automatic RECON loss notification
  - configuration 51
  - initialization 26
  - maintaining 143
  - parallel access 143
  - shutting down 143
- automatic restart management (ARM)
  - used in an IMSplex 21
- Automatic Restart Manager
  - element name 141
  - enabling 141
  - using 141
- Automatic Restart Manager (ARM) 79
- availability manager (AVM)
  - bringing up 608
  - definition of 600
  - operator messages 600
  - performing I/O prevention 600

- availability manager (AVM) (*continued*)
  - sending messages to operator 600
- AVM (availability manager)
  - bringing up 608
  - definition of 600
  - operator messages 600
  - performing I/O prevention 600
  - prefix for messages 600
  - sending messages to operator 600

## B

- back-end IMS
  - in a shared-queues environment 205
- backout
  - batch considerations 267
  - batch utility 267
  - definition 115
  - dynamic
    - data sharing 267
    - failure 267
  - overview 115
- backout with data sharing
  - dynamic 509
- backout, BACKOUT record for
  - recovery 548
- backup
  - database 111, 515
    - command for 516
  - message queue 112
  - overview 111
  - RECON data set 556
  - system data set 112
- BACKUP
  - definition 601
- backup copy
  - definition 111
- backup sessions for class-1 terminals
  - closing 597
  - logging on 597
  - opening 597
- backup sessions for XRF class-1 terminals
  - closing 591
  - opening 591
- BACKUP= keyword
  - defining class-1 terminals 636
- backward recovery
  - definition 115
- Base Primitive Environment (BPE)
  - as part of an IMSplex 21
  - commands 33
  - components that use BPE 33
  - configuration 33
  - defining an external trace data set 36
  - entry for z/OS PPT 127
  - messages 33
  - relationship between BPE and IMS 33
  - relationship to CSL 38
  - service provided 33
  - starting external tracing 38
  - stopping external tracing 38
  - tracing
    - enabling 35
  - user exit routines 33
  - writing to external data sets 36
  - writing to internal trace tables 35

- Base Primitive Environment (BPE) (*continued*)
  - z/OS PPT entry 127
- batch backout
  - DBRC support of 509
- batch backout utility
  - SSID naming convention 572
- Batch Backout utility (DFSBB000)
  - DBRC related 509
  - emergency restart, relation to 512
  - introduction to 115
  - restart, relation to 512
- batch environment
  - example 17
  - overview 17
  - TM 18
    - DB2 for z/OS, connecting 18
    - regions valid 18
- batch job
  - authorization 563
- batch jobs
  - tracking with DBRC 102
- batch jobs, converting to BMP jobs 279
- batch message processing (BMP) region
  - batch jobs, converting to BMP jobs 279
- Batch Terminal Simulator 364
- batch updater
  - scheduling 260
- binder return codes
  - interpreting 203
- BLDL list for dependent regions 412
- block size
  - OLDS 107
  - SLDS 110
- block-level sharing 233, 239
- BMP (batch message processing) region
  - application programs 8, 341
  - batch jobs, converting to BMP jobs 279
  - characteristics 8
  - DBCTL environment 73
  - Fast Path 73
- bottlenecks
  - message processing, detecting 828
- BPE (Base Primitive Environment)
  - commands 33
  - components that use BPE 33
  - configuration 33
  - defining an external trace data set 36
  - entry for z/OS PPT 127
  - List PROCLIB Member 168
  - messages 33
  - relationship between BPE and IMS 33
  - relationship to CSL 38
  - service provided 33
  - starting external tracing 38
  - stopping external tracing 38
  - tracing
    - enabling 35
  - user exit routines 33
  - writing to external data sets 36
  - writing to internal trace tables 35
  - z/OS PPT entry 127
- BTS (Batch Terminal Simulator)
  - highlights 364

- BTS (Batch Terminal Simulator) (*continued*)
  - simulating online execution 364
- buffer invalidation 271
- buffer pool
  - considerations 420
- buffer pools
  - allocating 416
  - analyzing requirements 416, 417
  - optimizing 416, 417
- buffers
  - OLDS 109
- business importance
  - Workload Manager 372

## C

- Call Summary report 362
- IMS Monitor
  - DB/DC 744
  - DCCTL 795
  - IMS Monitor (DBCTL) 776
- callable services
  - global 300
- calls from IMS to DBRC 481
- CANCEL command
  - non-XRF workload 632
- catalog management of data sets in RECON 574
- catch-up processing, definition 667
- CCTL regions 11
  - applications performance 414
  - contention for DRA resources 414
  - monitoring 389
  - performance objectives 391
  - PSB requests 414
- CCTL thread
  - abnormal termination 76
  - deadlock 76
  - Fast Path 73
  - scheduling a PSB 73
- central processor complex (CPC)
  - definition 586
- CFRM (coupling facility resource management)
  - couple data set format utility 188
  - policy, defining 179
- CFRM policy
  - defining 215
  - example 215
- CHANGE (/CHANGE) SURVEILLANCE command
  - XRF usage 597
- change accumulation
  - data set
    - naming convention 515
  - data sets
    - defining 105
    - reusing 105
  - definition of 489
  - group
    - defining 527
    - defining for future use 528
    - reusing 528
    - using 529
  - groups 105
  - logs 525



- change accumulation (*continued*)
  - nothing-to-process message 573
  - purge time 104
  - record
    - group 548
    - run 548
  - stops processing logs, during 573
  - utility 103
- change accumulation erroneously stops processing logs 573
- CHANGE.PRILOG command
  - changing RECON log control records 486
- CHANGE.RECON command
  - allocating space for RECON data sets 542
  - recovering RECONS 560
  - reorganizing RECON 559
- CHANGE.SECLOG command
  - changing RECON log control records 486
- changes in allocation and deallocation 497
- changing
  - system design 395
- changing applications
  - in an active IMS system 398
- channel-to-channel. 132
- checkpoint 91
  - concept 75
  - data set 180
  - for program synchronization 76
  - frequency 412
  - IMSplex, in an 304
  - monitoring processing effect
    - DB/DC 743
    - DBCTL 772
    - DCCTL 794
  - setting frequency 412
  - structure, initiating 181
  - system, initiating 180
- CHECKPOINT (/CHECKPOINT) command 75
- checkpoints
  - overhead of 91
  - system 91
- CIC parameter
  - for online image copy 519
- class-1 terminals
  - backup sessions opening 591
  - defining 636
  - defining priority of switching 648
  - definition of, for XRF 588
  - how takeover affects 588, 591
  - in XRF complex 634
  - ownership of in XRF 642
  - recommended ownership of 642
  - session recovery for XRF 597
  - specifying backup option 636
  - switching to backup sessions at post-takeover 621
  - switching to backup sessions at takeover 591, 616
  - takeover viewed by user 638
- class-2 terminals
  - defining priority of session recovery 648
- class-2 terminals (*continued*)
  - definition of, for XRF 588
  - establishing new sessions at post-takeover 621
  - establishing new sessions at takeover 616
  - how takeover affects 588, 616
  - in XRF complex 635
  - session recovery for XRF 597
  - takeover viewed by user 638
- class-3 terminals
  - definition of, for XRF 588
  - how takeover affects 588, 616
  - in XRF complex 636
- classification rule 373
- client
  - queue type 65
- cloned configuration
  - shared queues 211
- cloning
  - a shared-queues configuration 211
- CMD call
  - security overview 315
- CMDSEC= parameter 168
- COBOL subroutines, preloading 422
- cold start
  - global online change 463
  - local online change 463
  - multiple in a test environment 575
  - performing 463
  - when to perform 463
- COMCYCL parameter 141
- COMM macro statement
  - APPLID keyword 647, 648
  - PASSWORD keyword 647, 648
- command
  - CLEANUP.RECON
    - deleting unnecessary RECON records 558
  - DELETE.LOG
    - deleting unnecessary RECON records 558
  - INIT.CAGRP 527
- command authorization
  - DBRC, using DSPDCAX0 564
  - description 312
  - input command buffer 313
  - LU 6.2 313
  - using DSPDCAX0 and RACF 565
- command authorization exit routine
  - AO command security 320
- command security 157
  - command authorization exit routine 320
  - system definition 317
- commands
  - /CHECKPOINT 75
  - /DBDUMP database backup copies 517
  - /DISPLAY 441, 445
  - /ERESTART
    - restart after DBRC failure 512
    - restart after IMS failure 512
  - /MODIFY 441
  - /NRESTART, restart after IMS failure 512
  - /OPNDST 139
- commands (*continued*)
  - /RMGENJCL 493
  - /TRACE 378, 381
  - access method services (AMS)
    - REPRO 556
  - AO
    - security 315
  - authorizing connection to CQS
    - structures 178
  - authorizing CQS registration 178
  - BACKUP.RECON 556
  - CHANGE.PRILOG 486
  - CHANGE.RECON, recovering RECONS 560
  - CHANGE.SECLOG 486
  - creating an audit trail 153, 154
  - data sharing 250
  - DBCTL environment 11
  - DEFINE
    - GENERATIONDATAGROUP 36
  - DFSAPPL 120
  - effect on resources 351
  - for multiple resources 351
  - GENJCL 493
  - GENJCL.ARCHIVE 487
  - GENJCL.IC 531
  - GENJCL.OIC 531
  - global 170
  - INIT SELF 139
  - INIT.ADS 507
  - INIT.DB 507
  - INIT.DBDS 507, 521
  - INIT.RECON 536
  - INIT.RECON, establishing RECON data sets 505
  - INIT.RECON, recovering RECONS 560
  - INITIATE OLC 445
    - without RM (resource manager) 445
  - log-related
    - CHANGE.PRILOG 488
    - CHANGE.RECON 488
    - CHANGE.SECLOG 488
    - DELETE.LOG 488
    - GENJCL.ARCHIVE 488
    - GENJCL.CLOSE 488
    - LIST.LOG 488
    - NOTIFY.PRILOG 488
    - NOTIFY.SECLOG 488
  - NOTIFY.IC
    - HALDB online reorganization considerations 531
  - NOTIFY.UIC 515
    - HALDB online reorganization considerations 531
  - processing considerations in a CSL 151
  - QUERY 445
  - QUERY OTMATI 383
  - REPRO 560
  - requirements 68
  - restricting for RECON data set 563
  - routing 151
  - TERMINATE 445, 456
  - TRACE CT 380
  - type-2 38

- commands (*continued*)
  - configuration requirements 68
  - environment 68
  - type-2 command environment
    - restrictions for online change without RM 49
  - VTAM VARY 139
- COMMIT command 442, 446
- commit phase 1
  - failure
    - scenario 460
- commit point 91
  - application program 92
- Common Queue Server (CQS)
  - access with IMS commands 223
  - address space 38
  - as part of an IMSplex 21
  - cold start 219
  - coupling facility
    - changing structure size 219
  - definitions
    - execution parameters 218
    - global structure parameters 218
    - initialization parameters 218
    - local structure parameters 218
  - entry for z/OS PPT 127
  - list structures
    - planning 224
  - logging 67
    - OC/390 logger 219
  - planning
    - hardware 62
    - software 62
  - restarting 219
    - after system checkpoint 219
  - starting 219
  - structure recovery 66
  - structures
    - checkpoint, initiating 219
    - copying 219
    - deleting 219
    - rebuilding 219
    - recovering 219
  - system checkpoint
    - data sets 219
    - initiating 219
    - restart 219
  - user-supplied exit routines 219
  - using 219
  - warm start 219
    - log token 219
  - z/OS PPT entry 127
- Common Service Layer (CSL)
  - address spaces 41
  - administering 141
  - as part of an IMSplex 21
  - automatic RECON loss
    - notification 51
  - benefits 38
  - command processing
    - considerations 151
  - configuration recommendation 54
  - configured without Resource Manager 41
  - failure 705
  - IMS address spaces included 38
  - introduction 38
- Common Service Layer (CSL) (*continued*)
  - managers 41
  - minimum configuration 54
  - ODBM (Open Database Manager)
    - overview 41
  - Open Database Manager (ODBM)
    - overview 41
  - Operations Manager
    - overview 44
  - overview 38
  - Resource Manager
    - function provided 44
    - overview 44
    - resource structure 44
  - Structured Call Interface
    - functions provided 45
    - overview 45
- Communication Controller
  - 37x5 595
- Communication IWAIT report 428, 430
- Communication Summary report 428
  - IMS Monitor (DB/DC) 754
  - IMS Monitor (DCCTL) 803
- Communication Wait report
  - IMS Monitor (DB/DC) 755
  - IMS Monitor (DCCTL) 804
- component, as part of RSR name 669
- components of a CQS 62
- compression
  - PRILOG 557
- concurrent image copy (CIC)
  - database backup copies 519
  - registered database with DBRC 519
  - restrictions 519
- configurations
  - IMSplex
    - DBCTL 54
    - minimum 54
    - mixed IMS versions 54
    - recommendations 54
    - shared queues with a CSL 54
    - shared queues without a CSL 54
    - single system 54
    - typical 54
    - without Resource Manager 54
  - IMSplex with multiple SPOC
    - users 45
  - simple 41
- connection failures
  - coupling facility 283
- considerations
  - system administration 125
- contention
  - avoiding with RECON data set 536
- control program 8
- control region
  - address spaces
    - allocating real storage 427
  - control program 8
  - DCCTL environment, in a 13
  - definition 8
  - DL/I calls 8
  - execution parameters
    - defining in a shared-queues environment 218
  - IMSplex, in a 24
  - restarting 355
- control region (*continued*)
  - starting 355
- CONTROLINTERVALSIZE keyword
  - DEFINE CLUSTER keywords 542
- conversation mode
  - IMSplex, in an 304
- conversational transactions
  - in a shared-queues environment 211
- conversion
  - batch jobs to BMP jobs 279
- copy
  - structures 188
- corrective service
  - installing 195
  - IMSplex 200
- couple data set format utility 188
- coupling facilities
  - monitoring structures 352
- coupling facility
  - as part of an IMSplex 21
  - connection failures 283
  - defining CFRM policy
    - example 215
  - failure due to IMS restart 286
  - failure in rebuilding 286
  - monitoring structures 254
  - resource structure 53
  - structure failures 285
  - structures
    - changing size 283
    - structures, calculating size of 280
    - sysplex data sharing 273
- coupling facility resource management (CFRM)
  - couple data set format utility 188
  - policy, defining 179
- CPC (central processor complex)
  - definition 586
  - failure
    - as cause of XRF takeover 591, 613
    - requirement for XRF 588, 595
- CPI-C driven application programs
  - security considerations in 321
- CQS (Common Queue Server)
  - access with IMS commands 223
  - address space 38
  - administration 177
  - authorization 178
  - benefits 62
    - automatic work load balancing 62
    - incremental growth 62
    - reliability 62
  - client connection
    - establishing 177
  - client failure 177
  - cold start 219
  - components 62
    - checkpoint data set 62
    - overflow structure 62
    - primary structure 62
    - resource structure 62
    - structure recovery data set 62
    - z/OS log stream 62
  - coupling facility
    - changing structure size 219
  - defining 24



## CQS (Common Queue Server) *(continued)*

- definitions
    - execution parameters 218
    - global structure parameters 218
    - initialization parameters 218
    - local structure parameters 218
  - diagram of client systems and coupling facility 62
  - entry for z/OS PPT 127
  - functions
    - overflow processing 62
    - records restart 62
    - requests 62
    - structure checkpoint 62
    - structure rebuild 62
    - system checkpoint 62
  - information for restart 177
  - list structures
    - planning 224
  - logging 67
    - OC/390 logger 219
  - notification of work 62
  - operating system requirements 62
  - overview 62
  - planning
    - hardware 62
    - software 62
  - rebuilding structures 185
  - recovering 185
  - resource structure
    - attributes 53
  - restarting 219
    - after system checkpoint 219
  - restarting information 177
  - restating after system checkpoint 180
  - starting 219
  - structure overflow function 183
  - structure recovery 66
  - structure recovery data set 99
  - structures
    - checkpoint, initiating 219
    - copying 219
    - deleting 219
    - rebuilding 219
    - recovering 219
  - system checkpoint 219
    - data sets 219
    - initiating 219
    - restart 219
  - system checkpoint data set 98
  - user-supplied exit routines 219
  - using 219
  - warm start 219
    - log token 219
  - z/OS log data set 98
  - z/OS PPT entry 127
- ## CQS benefits 62
- ## CQS client
- requests 67
- ## CQS-managed rebuild 186
- ## cross-system extended services (XES)
- used in an IMSplex 21
- ## cross-system queuing, determining 764
- ## cryptography 341
- ## CSBLK DSECT 300
- ## CSL (Common Service Layer) 19
- address spaces 41

## CSL (Common Service Layer) *(continued)*

- administering 141
  - benefits 19, 38
  - configuration recommendation 54
  - configuration rules 52
  - configured without Resource Manager 41
  - CSLZQRY 144
  - failure 705
  - IMS address spaces included 38
  - introduction 38
  - managers 41
  - minimum configuration 54
  - overview 38
  - relationship to Base Primitive Environment 38
  - Resource Manager
    - function provided 44
    - overview 44
    - resource structure 44
  - security 30
  - Structured Call Interface
    - functions provided 45
    - overview 45
  - XRF 585
- ## CSLDCxxx
- ODBM configuration 24
- ## CSLDIxxx
- ODBM initialization 24
- ## CSLOIxxx 25
- ## CSLSIxxx 25
- ## CSLSREG 175
- ## CTC (channel-to-channel)
- failure, XRF limitations 591
- ## CTC (channel-to-channel) channel-end
- appendage 132
- ## CTRACE records 242
- ## CTRACE service 380
- ## CYLINDERS keyword
- DEFINE CLUSTER keywords 542

## D

- ## damaged RECON data set 560
- ## data
- how applications access 237
- ## Data Capture exit routine 110
- ## Data Capture exit routines
- description 78
- ## data dictionary
- IBM DB/DC Data Dictionary 84
- ## data entry database 72
- ## data entry database (DEDB)
- records in the RECON data set 549
  - resource information stored 172
- ## data propagation 78
- ## data set allocation
- RECON data set
    - creating 542
    - shared among multiple processors 538
- ## data set placement
- XRF 656
- ## data set, RECON 535
- ## data sets
- A-through-J 531
  - ISPTABL 120

## data sets *(continued)*

- M-through-V 531
- ## data sharing
- area-level 233, 234
  - assigning a sharing level with DBRC 499
  - backout 267
  - block level 499
  - block-level 239
  - commands 250
  - configuration differences 233
  - configuration examples 245
  - controlling 352
- ## database
- reorganizing 259
- ## database as a data resource 233
- ## database availability 233
- ## database integrity 236, 499
- ## database level 499
- multiple readers 244
  - update activity 243
- ## database processing
- starting 250
  - stopping 250
- ## database share level
- changing 264
- ## database-level 233, 234
- ## DB/DC environment 10
- ## DBCTL environment, in a 12
- ## DBRC 257, 268, 352
- commands 258
- ## DBRC controlled 499
- ## DBRC failure
- restart after 269
- ## DBRC, role of 234
- ## DEDB availability 233
- ## description 233, 239
- ## dynamic backout 509
- ## establishing naming conventions 236
- ## excluding 248
- ## forward recovery 268
- ## image copies 237
- ## image copy 260, 263
- making 263
- ## IMS failure 288
- restart after 268
- ## information in the RECON data set 499
- ## installation tasks 246
- ## IRLM
- restarting 269
  - starting 249
  - stopping 250
- ## IRLM, role of 239
- ## levels of 499
- ## logs 265
- ## merging logs 526
- ## normal operations 260
- ## online change 259
- ## planning for recovery 509
- ## procedures 266
- ## RECON, monitoring 254
- ## record 554
- ## recovery 265
- ## resources, monitoring 252, 352
- ## RSR support 682
- ## starting 249

- data sharing (*continued*)
  - stopping 249
  - sysplex 270, 283
    - concepts 271
    - defining the group 243
    - defining the lock structure 243
    - defining the users 243
    - IRLM 243
    - terminology 271
- data space
  - RACF 322
- database
  - access intent, changing 257
  - allocation record 554
  - application design review 82
  - authorization, changing 572
  - availability 259
  - backup 111, 515
    - commands for 516
  - backup guidelines 524
  - damage from backout processes 509
  - deallocation
    - logging in the RECON data set 497
  - DEDB 72
  - dynamic recovery
    - non-data sharing 509
    - with data sharing 509
    - without data sharing 509
  - Fast Path types 72
  - forward recovery
    - dynamic non-data sharing 509
    - dynamic with data sharing 509
    - non-data sharing 509
    - with data sharing 509
    - without data sharing 509
  - groups 549
  - I/O analysis 434
  - image copy, making 263
  - integrity in data sharing 499
  - making HISAM copies 523
  - making image copies 516
  - monitoring buffers 753, 780
  - monitoring DL/I calls 741, 791
  - MSDB 72
  - page fixing buffers 416
  - processing during takeover 614
  - productivity tools 350
  - protection 339
  - record 549
  - recovery
    - batch 509
    - dynamic non-data sharing 509
    - dynamic with data sharing 509
  - recovery records 548
  - reorganizing 259, 262
  - requirement for XRF 588
  - shave level change 264
- database allocation record 554
- database buffer pool
  - optimization 419
- Database Buffer Pool report
  - IMS Monitor (DBCTL) 780
- Database Change Accumulation utility (DFSUCUM0)
  - CA group
    - defining 527
- Database Change Accumulation utility (DFSUCUM0) (*continued*)
  - CA group (*continued*)
    - defining for future use 528
    - not reusing 528
    - reusing 528
  - condensing SLDS or RLDS 102
  - description 525
  - execution recorded by DBRC 527
  - input 103
  - overview 525
  - purge time 104
  - subset of log volumes 526
  - usage guidance 526
  - valid log subset with DBRC 526
- database data set
  - initiating in the RECON data set 507
- database data set (DBDS)
  - A-through-J, L and X 531
  - M-through-V and Y 531
- database data set (DBDS) group
  - commands that affect the definition
    - CHANGE.DBDSGRP 529
    - DELETE.DBDS 529
    - DELETE.DBDSGRP 529
    - INIT.DBDSGRP 529
  - commands that specify
    - GENJCL.IC 529
    - GENJCL.OIC 529
    - GENJCL.RECEIVE 529
    - GENJCL.RECOV 529
    - GENJCL.USER 529
    - LIST.DBDS 529
    - LIST.HISTORY 529
  - ILDS (Indirect List Data Set) 529
  - using 529
- database data set record 549
- Database Image Copy 2 utility (DFSUDMT0)
  - overview 518
- Database Image Copy utility (DFSUDMP0)
  - creating data sets for future use 520
  - description 516
  - execution recorded by DBRC 515
  - image copies, creating 517
  - maximum number of generations 521
  - nonstandard image copy data sets 524
  - recovery period 521
  - reusing image copy data sets 522
- database integrity
  - data sharing support 236
- DATABASE macro statement
  - data sharing, and 238
  - used for online change 403
- database readiness level tracking (DLT) 672
- database recovery
  - Database Image Copy 2 output
    - format 519
  - database recovery
    - concepts 489
  - DBRC role 493
  - groups 549
  - process overview 489
- Database Recovery Control 361
- database recovery control (DBRC)
  - as part of an IMSplex 21
- Database Recovery Control (DBRC)
  - active RECON data set 544
  - API request authorization
    - using a security product 563
  - API requests
    - overview 482
  - assigning a sharing level 499
  - authorization
    - changing database 572
  - batch commands
    - BACKUP.RECON command 482
    - CHANGE command 482
    - CLEANUP.RECON command 482
    - DELETE command 482
    - GENJCL command 482
    - INIT command 482
    - LIST command 482
    - NOTIFY command 482
    - RESET.GSG command 482
  - calls from IMS 481
  - catalog management of data sets 574
  - change accumulation
    - nothing-to-process message 573
  - closing an open online PRILOG record 570
  - command authorization 579
    - using a security product 563
    - using both DSPDCAX0 and RACF 565
    - using DSPDCAX0 564
  - commands
    - issuing 480
  - communicating with 482
  - considerations for using 513
  - data set catalog management 574
  - data set naming conventions 513
  - data sets
    - partitioned 480
    - RECON 480
    - RECON data set 535
  - data sharing control 499
  - data-sharing records 499
  - database backup
    - Concurrent Image Copy 519
    - controlling the number of image copies managed 521
    - Database Image Copy (DFSUDMP0) 517
    - Database Image Copy 2 (DFSUDMT0) 518
    - guidelines 524
    - HISAM copies 523
    - methods 515
    - Online Database Image Copy (DFSUICP0) 519
    - recovering a database 493
    - recovery period of image copy data sets and GENMAX 521
  - Database Recovery Control utility (DSPURX00) 480
  - database recovery for RSR 501
  - DBRC parameter
    - IMSCRTL macro 505
  - DBRC procedure 505

Database Recovery Control (DBRC)  
(continued)

- generating JCL 480
- GENMAX parameter
  - using 521
- GENMAX, resetting 569
- groups
  - definition of 530
- IMS procedures and 505
- IMSpIex support 503
- initial RECON access 544
- initialization 505
- initializing the RECON data set 505
- introduction 8
- IRLM status 512
- log control 485
- log data sets
  - notifying of moves 574
- log-related commands
  - CHANGE.PRILOG 488
  - CHANGE.RECON 488
  - CHANGE.SECLOG 488
  - DELETE.LOG 488
  - GENJCL.ARCHIVE 488
  - GENJCL.CLOSE 488
  - LIST.LOG 488
  - NOTIFY.PRILOG 488
  - NOTIFY.SECLOG 488
- OLR, considerations for 531
- online commands
  - /RMCHANGE command 482
  - /RMDELETE command 482
  - /RMGENJCL command 482
  - /RMINIT command 482
  - /RMLIST command 482
  - /RMNOTIFY 482
- overview 479
- parameters
  - IMS.PROCLIB execution
    - parameter 505
- partitioned data set members 480
- PRILOG record
  - closing 570
- RECON data set
  - accessing in parallel 536
  - ALLOC record 554
  - avoiding contention 536
  - avoiding deadlock situations 538
  - backing up 556
  - backing up before
    - reorganization 559
  - BACKOUT record 548
  - CA record 548
  - CAGRP record 548
  - command authorization
    - support 563
  - contention, avoiding 536
  - creating 542
  - DBDSGRP record 549
  - DEFINE CLUSTER keywords 542
  - extending 542
  - GSG record 552
  - header records 546
  - IMAGE record 553
  - initializing 536
  - input/output error
    - processing 560

Database Recovery Control (DBRC)  
(continued)

- RECON data set (continued)
  - log data set records 546
  - LOGALL record 554
  - loss notification 562
  - maintaining 556
  - parallel access 536
  - parallel access, planning for 537
  - planning considerations 535
  - processing using LSR option 542
  - recommendations when
    - creating 542
  - record mapping 555
  - record types 545
  - RECOV record 554
  - recovering 560
  - recovery record types 548
  - registering databases in 507
  - registering DBDSs in 507
  - REORG record 553
  - reorganizing online 559
  - reorganizing overview 559
  - reorganizing procedure 559
  - replacing a discarded
    - RECON 562
  - replacing damaged 560
  - request authorization support 563
  - security considerations 544
  - SSYS record 554
- RECON initialization token (RIT) 497
- recording image copies 515
- recovery involving IRLM 512
- recovery record types 548
- recovery utilities 494
- RECOVPD parameter
  - using 521
- registering databases in 507
- registering DBDSs in 507
- resource names for authorization 579
- RIT (RECON initialization token) 497
- security resource profiles 563
- sharing level, assigning 499
- skeletal JCL 480
- SLDS stop time, locating the last in
  - the RECON data set 567
- spare RECON data set 544
- specifying when it is to be used 505
- subsystem (SSYS) records 572
- system considerations 513
- tasks performed with 481
- usage tips 567
- validating utility JCL 494
- when to use 479

database tracker

- DL/I 668
- Fast Path 669
- milestone, definition 668

Database-Buffer-Pool report

- description 720
- fields in the report 720
- IMS Monitor (DB/DC) 753
- using the report 720

database-level data sharing 233, 234

databases

- accessing 353
- dynamically allocation in XRF 656

databases (continued)

- ensuring integrity 600
- ensuring integrity with XRF 591
- failure, XRF limitations 591
- resource information stored 171

databases supported

- DB/DC environment 10

DB batch 17

DB Groups

- commands that affect
  - CHANGE.DBDSGRP 530
  - DELETE.DBDSGRP 530
  - INIT.DBDSGRP 530
  - LIST.DBDSGRP 530

DB header record

- HALDB 549

DB Monitor 378

DB Monitor reports

- Database-Buffer-Pool report 720
- Distribution-Appendix report 727
- DL/I-Call-Summary report 724
- Monitor-Overhead report 730
- Program-I/O report 722
- VSAM-Buffer-Pool report 709
- VSAM-Statistics report 715

DB partition record

- HALDB 549

DB/DC

- security
  - overview 307

DB/DC Data Dictionary

- network documentation 84
- system documentation 84

DB/DC environment

- communication design 402
- control region 8
- data sharing 10
- databases supported 10
- definition 6
- example 6
- RSR 11
- XRF 10

DB2 for z/OS

- Data Capture exit routines 78
- recovery using RSR 688

DB2 Recoverable Resource Manager

- Services Attach Facility (DB2RRMS)
  - using in JMP or JBP applications 475

DB2 Recoverable Resource Services

- attachment facility (RRSAF) 13

DB2RRMS (DB2 Recoverable Resource Manager Services Attach Facility)

- using in JMP or JBP applications 475

DBBBATCH

- IMSpIex component 39

DBCTL

- security facilities 342

DBCTL (Database Control)

- environment
  - Fast Path 73
- overview 113
- recovery 113
- XRF capabilities 591

DBCTL (Database Control) environment

- monitoring 389
- security 341

- DBCTL environment
  - CCTL 11
  - commands 11
  - control 11
  - data sharing 12
  - databases supported 12
  - DL/I 11
  - DRA 11
  - overview 11
  - terminals 11
  - testing 363
  - XRF alternate 13
- DBCTL standby emergency restart 449
- DBDS
  - initiating in the RECON data set 507
- DBDS (database data set)
  - A-through-J, L and X 531
  - M-through-V and Y 531
- DBDS (database data set) group
  - commands that affect the definition
    - CHANGE.DBDSGRP 529
    - DELETE.DBDS 529
    - DELETE.DBDSGRP 529
    - INIT.DBDSGRP 529
  - commands that specify
    - GENJCL.IC 529
    - GENJCL.OIC 529
    - GENJCL.RECEIVE 529
    - GENJCL.RECOV 529
    - GENJCL.USER 529
    - LIST.DBDS 529
    - LIST.HISTORY 529
  - ILDS (Indirect List Data Set) 529
  - record 549
  - using 529
- DBDUMP (/DBDUMP) command
  - database backup copies 517
- DBFULTA0 (Fast Path Log Analysis utility) 349
- DBRC
  - administration 479
  - security 563
- DBRC (Database Recovery Control)
  - active RECON data set 544
  - API request authorization
    - using a security product 563
  - API requests
    - overview 482
  - as part of an IMSplex 21
  - assigning a sharing level 499
  - authorization
    - changing database 572
  - automatic RECON loss
    - notification 51
  - batch commands
    - BACKUP.RECON command 482
    - CHANGE command 482
    - DELETE command 482
    - GENJCL command 482
    - INIT command 482
    - LIST command 482
    - NOTIFY command 482
    - RESET.GSG command 482
  - calls from IMS 481
  - catalog management of data sets 574
  - change accumulation erroneously
    - stops processing logs 573

- DBRC (Database Recovery Control)
  - (continued)
  - change accumulation
    - nothing-to-process message 573
  - closing an open online PRILOG
    - record 570
  - command authorization 579
    - using a security product 563
    - using both DSPDCAX0 and RACF 565
    - using DSPDCAX0 564
  - commands
    - INIT.CA 528
    - INIT.CAGRP 528
    - issuing 480
  - communicating with 482
  - components 480
  - considerations for using 513
  - controlling database recovery 489
  - data set catalog management 574
  - data set naming conventions 513
  - data sets
    - defining recovery
      - requirements 489
      - partitioned 480
    - RECON 480, 556
    - RECON data set 535
    - SSYS record 554
  - data sharing 257, 268, 352
    - commands 258
  - data sharing control 499
  - data sharing control, and 233, 234
  - data-sharing records 499
  - database backup
    - Concurrent Image Copy 519
      - controlling the number of image copies managed 521
    - creating image copy data sets for future use 520
  - Database Image Copy
    - (DFSUDMP0) 517
  - Database Image Copy 2
    - (DFSUDMT0) 518
  - guidelines 524
  - HISAM copies 523
  - methods 515
  - Online Database Image Copy
    - (DFSUICP0) 519
  - recovering a database 493
  - recovery period of image copy
    - data sets and GENMAX 521
  - reusing image copy data sets 522
- Database Recovery Control utility
  - (DSPURX00) 480
- database recovery for RSR 501
- DBRC address space, initiating 8
- DBRC components 480
- DBRC parameter
  - IMSCtrl macro 505
- DBRC procedure 505
- DSPSCIX0 26
- generating JCL 480
- GENMAX parameter
  - using 521
- GENMAX, resetting 569
- groups
  - definition of 530

- DBRC (Database Recovery Control)
  - (continued)
  - IMS procedures and 505
  - IMSplex support 503
  - initial RECON access 544
  - initialization 505
  - initializing the RECON data set 505
  - introduction 8, 113
  - IRLM status 512
  - log control 485
  - log data sets
    - notifying of moves 574
  - log records, deleting 556
  - log-related commands
    - CHANGE.PRILOG 488
    - CHANGE.RECON 488
    - CHANGE.SECLOG 488
    - DELETE.LOG 488
    - GENJCL.ARCHIVE 488
    - GENJCL.CLOSE 488
    - LIST.LOG 488
    - NOTIFY.PRILOG 488
    - NOTIFY.SECLOG 488
  - maintaining RECON records 556
  - OLR, considerations for 531
  - online commands
    - /RMCHANGE command 482
    - /RMDELETE command 482
    - /RMGENJCL command 482
    - /RMINIT command 482
    - /RMLIST command 482
    - /RMNOTIFY 482
  - overview 479
  - parameters
    - IMS.PROCLIB execution
      - parameter 505
  - partitioned data set members 480
  - PRILOG compression 557
  - PRILOG record
    - closing 570
  - RECON data set 98
    - accessing in parallel 536
    - ALLOC record 554
    - avoiding contention 536
    - avoiding deadlock situations 538
    - backing up 556
    - backing up before
      - reorganization 559
    - BACKOUT record 548
    - CA record 548
    - CAGRP record 548
    - command authorization
      - support 563
    - contention, avoiding 536
    - creating 542
    - DBDSGRP record 549
    - DEFINE CLUSTER keywords 542
    - extending 542
    - GSG record 552
    - header records 546
    - IMAGE record 553
    - initializing 536
    - input/output error
      - processing 560
    - log data set records 546
    - LOGALL record 554
    - loss notification 562

- DBRC (Database Recovery Control)
  - (continued)
  - RECON data set (continued)
    - maintaining 556
    - parallel access 536
    - parallel access, planning for 537
    - planning considerations 535
    - processing using LSR option 542
    - recommendations when
      - creating 542
    - record mapping 555
    - record types 545
    - RECOV record 554
    - recovering 560
    - recovery record types 548
    - registering databases in 507
    - registering DBDSs in 507
    - REORG record 553
    - reorganizing online 559
    - reorganizing overview 559
    - reorganizing procedure 559
    - replacing a discarded
      - RECON 562
    - replacing damaged 560
    - request authorization support 563
    - security considerations 544
    - SSYS record 554
  - RECON data set, defining recovery
    - requirements 489
  - RECON initialization token (RIT) 497
  - recording change accumulations 527
  - recording image copies 515
  - recovery involving IRLM 512
  - recovery record types 548
  - recovery utilities 494
  - RECOVPD parameter
    - using 521
  - registering databases in 507
  - registering DBDSs in 507
  - required JCL updates 248
  - resource names for authorization 579
  - restart 269
  - RIT (RECON initialization token) 497
  - security 563
  - security resource profiles 563
  - sharing level, assigning 499
  - skeletal JCL 480
  - SLDS stop time, locating the last in
    - the RECON data set 567
  - spare RECON data set 544
  - specifying when it is to be used 505
  - subsystem (SSYS) records 572
  - system considerations 513
  - tasks performed with 481
  - testing 361
  - usage tips 567
  - validating utility JCL 494
  - when to use 479
- DBRC Command Authorization exit routine (DSPDCAX0)
  - usage 564
- DBRC Command Authorization Support
  - using both DSPDCAX0 and RACF 565
- DBRC SCI Registration exit routine (DSPSCIX0)
  - defining IMSplex name 26
- DBRC security
  - overriding security
    - RECON data set 565
- DCCTL
  - security
    - overview 307
- DCCTL environment
  - AO transactions 16
  - application calls supported 16
  - application programs 13
  - communication design 402
  - connecting subsystems 13
  - control region
    - data communication manager 13
    - message manager 13
    - Transaction Manager 13
  - data communication calls 16
  - databases supported 15
  - DBRC 13
  - definition 16
  - diagnosis 16
  - ESAF 13, 15
  - example 13
  - exit routines 16
  - external subsystems 13
  - IMS.PROCLIB 15
  - initialization 16
  - overview 13
  - region types 13
  - restart 16
  - RRSAF 15
  - stage 1 input 16
  - status code AD 16
  - system service calls 16
  - termination 16
  - testing 364
- DDM (distributed data management)
  - Open Database Manager (ODBM)
    - administration 147
    - overview 41
- deadlock
  - abnormal termination 76
  - CCTL thread 76
  - DBCTL 76
  - possible causes 76
- Deadlock Event Summary report
  - IMS Monitor (DB/DC) 755
  - IMS Monitor (DBCTL) 781
- deadlock, avoiding RECON data set 538
- deallocation
  - logging in the RECON data set 497
- DEDB (data entry database)
  - access intent, changing 257
  - DBCTL environment 73
  - Fast Path considerations 72
  - online change 443, 447
  - records in the RECON data set 549
- default security 307
  - definition 308
  - security option 308
- DEFINE CLUSTER
  - keyword
    - recommended values 542
- DEFINE CLUSTER keywords
  - CONTROLINTERVALSIZE 542
  - CYLINDERS 542
  - FREESPACE 542
- DEFINE CLUSTER keywords (continued)
  - INDEXED 542
  - KEYS 542
  - LOG 542
  - NAME 542
  - NOWRITECHECK 542
  - RECORDSIZE 542
  - SHAREOPTIONS 542
  - SPEED 542
- DEFINE GENERATIONDATAGROUP
  - command 36
- degraded mode logging 106
- DELAY parameter 141
- deleting a SSYS record 572
- deleting information
  - log records from RECON 558
- deleting log records, how to 556
- dependent region
  - address spaces
    - allocating real storage 427
- dependent regions
  - BLDL list 412
  - definition 8
  - security 327, 341
  - resource access security (RAS) 328
  - security options for 70
  - starting 8
  - starting on alternate IMS system 610
  - status recorded in log 611
  - types of 8
- dependent service element (DSE)
  - definition 586
- DEQ macro 538
- deregistering interest
  - in transactions 207
  - shared queues
    - in LTERMs 207
    - in MSC (multiple systems coupling) resources 208
- design reviews 82
  - participation 82
- Destination Creation exit routine (DFSINSX0)
  - shared queues
    - dynamic control blocks 209
    - message switching 209
    - program-to-program switching 209
    - undefined transactions
      - dynamic control blocks 209
- detecting bottlenecks in message processing 828
- DFP Record Management Services 538
- DFSAPPL command 120
- DFSBB000 (Batch Backout utility)
  - emergency restart, relation to 512
  - introduction to 115
  - limitations 509
  - restart, relation to 512
- DFSCCBLK 300
- DFSCGxxx
  - global online change 26



- DFSCGxxx (*continued*)
  - OLCSTAT data set 26
- DFSCMC10 module
  - installing IMS 129
  - installing MSC CTC links 129
- DFSCSGN0 module for signon verification 331
- DFSCSTRN0 module for transaction authorization 331
- DFSDCxxx PROCLIB member
  - signon verification 311
- DFSERA10 (File Select and Formatting Print utility) 349
- DFSFIXxx member of IMS.PROCLIB
  - DFSINTRS, used to page fix intent lists 420
  - for page fixing 420
  - used in tuning 420
- DFSHSBxx member of IMS.PROCLIB
  - choosing methods of XRF surveillance 597
  - establish criteria for takeover 613
- DFSILTA0 (Log Transaction Analysis utility) 349
  - used for monitoring 375
- DFSINSX0 (Destination Creation exit routine)
  - shared queues
    - dynamic control blocks 209
    - message switching 209
    - program-to-program switching 209
- DFSINTRS, used to page fix intent lists 416
- DFSINTX0 (Initialization exit routine)
  - using with ETO 327
- DFSIRP0 (Program Isolation Trace Report utility) 381
- DFSISTS0 (Statistical Analysis utility) 349
- DFSKBLA3 (Knowledge-Based Basic Formatting Print routine) 349
  - overview of 833
- DFSKBLAK (Knowledge-Based Formatting Print routine) 349
  - overview of 833
- DFSKBLAS (Knowledge-Based Summary Formatting Print routine) 349
  - overview of 833
- DFSLGNX0 (Logon exit routine)
  - using with ETO 327
- DFSMPLxx, for program preload 422
- DFSMRCL0 (resource cleanup module)
  - uninstalling 133
- DFSMS
  - requirement for XRF 588, 595
  - XRF process, contribution to 601
- DFSOFMD0
  - offline dump formatting module 136
- DFSPPREC0 (Index/ILDS Rebuild utility)
  - relation to DBRC 494
- DFSSGNX0 (Sign-on exit routine)
  - using with ETO 327
- DFSUARC0 (Log Archive utility) 95, 96
  - description 487
  - OLDS archive 99
- DFSUCUM0 (Database Change Accumulation utility) 102
  - CA group
    - defining 527
    - defining for future use 528
    - not reusing 528
    - reusing 528
  - description 525
  - execution recorded by DBRC 527
  - overview 525
  - subset of log volumes 526
  - usage guidance 526
- DFSUDMP0 (Database Image Copy utility)
  - creating data sets for future use 520
  - description 516
  - maximum number of generations 521
  - nonstandard image copy data sets 524
  - recovery period 521
- DFSUDMT0 (Database Image Copy 2 utility)
  - overview 518
- DFSUIC00 260
- DFSUICP0 (Online Database Image Copy utility)
  - creating data sets for future use 520
  - description 516
  - execution recorded by DBRC 515
- DFSUOLC0 463
- DFSUOLC0 (Global Online Change utility) 449, 451
- DFSURUL0
  - HISAM Reorganization Unload utility 111
  - introduction to 111
- DFSURUL0 (HISAM Reorganization Unload utility)
  - for backup 523
- DFSVMxx, for buffer subpools
  - used in tuning 420
  - using IOBF 420
  - using VSAMFIX on OPTIONS statement 420
- dial-up lines 642
- dictionary
  - IBM DB/DC Data Dictionary 84
  - using a data dictionary 84
- differences in takeover process 618
- disabling enforcement
  - resource name uniqueness 299
  - resource type consistency 300
- discarded RECON, replacing 562
- dispatching priority 428
- DISPLAY (/DISPLAY) command
  - for monitoring 375
  - MODIFY 441, 445
  - use in accounting 832
- display bypass security 338
- display for structure full threshold 183
- distributed data management (DDM)
  - Open Database Manager (ODBM)
    - administration 147
    - overview 41
- Distributed Relational Database Access (DRDA)
  - Open Database Manager (ODBM)
    - administration 147
    - overview 41
- Distribution Appendix report
  - IMS Monitor (DB/DC) 760
  - IMS Monitor (DBCTL) 783
  - IMS Monitor (DCCTL) 806
- Distribution-Appendix report
  - description 727
  - generating the report 729
  - interpreting 762
  - interpreting output 786
  - output 786
  - redefining default ranges 730
- DL/I database tracker 668
- DL/I databases
  - considerations 425
  - design considerations 425
- DL/I-Call-Summary report
  - description 724
  - fields in the report 725
  - using the report 724
- DLIBATCH
  - IMSplex component 39
- DLISAS procedure, security 339
- DMB buffer pool
  - analyzing requirements 418
- documentation
  - network 85
  - production configuration 85
  - system definition 85
  - terminal profiles 85
- documenting
  - data dictionary 84
  - IBM DB/DC Data Dictionary 84
  - the IMS system 85
- documenting IMS system 81
- DRA (database resource adapter)
  - overview 11
  - resources
    - CCTL transaction requests 389
    - exceeded by PSB schedule requests 414
- DRD
  - RSR 684
  - support on RSR 684
- DRD (dynamic resource definition)
  - IMSplex scenarios 29
- DRDA (Distributed Relational Database Access)
  - Open Database Manager (ODBM)
    - administration 147
    - overview 41
- DSE (dependent service element)
  - definition 586
- DSPDBHRC 549
  - partition DB record
    - HALDB 549
- DSPDCAX0 (DBRC Command Authorization exit routine)
  - usage 564
- DSPDSHRC
  - partition DBDS records
    - HALDB 549

- DSPTNRC
  - HALDB partition record
  - HALDB 549
- dual logging 106
  - definition 94
- dump format 519
- duplexing
  - explicitly stopping 188
  - structure 66, 188
  - unnecessary overhead 188
- dynamic allocation
  - ACB library 32
  - area data sets in XRF 656
  - availability under z/OS 32
  - database data sets 32
  - HALDBs 32
  - IMS databases in XRF 656
  - OLDS 32
  - RECON data set 541
  - RECON data sets 32
  - SLDS 32
  - WADS 32
- dynamic bailout
  - data sharing 509
  - non-data sharing 509
- dynamic database bailout
  - data sharing 509
  - non-data sharing 509
- dynamic resource definition
  - IMS resources
    - create 447
    - delete 447
    - update 447
- dynamic resource definition (DRD)
  - IMSplex scenarios 29

## E

- E-MCS 45
- EEQE (Extended Error Queue Element),
  - in XRF complex 621
- element names, ARM 141
- EMH (Expedited Message Handler)
  - shared queues environment, in a 229
- enabling
  - shared queues 215
- encryption
  - using Segment Edit/Compression
    - exit 341
  - VTAM terminals 341
- end-user service
  - at takeover 588
  - at XRF takeover 591
- ENF 186
- enqueue problems, causes of
  - RECON 576
- enqueue/dequeue tables
  - use 76
- ERESTART (/ERESTART) BACKUP
  - command 608, 621
- error
  - application logic 118
  - INITIATE OLC command 454
  - input 118
  - operational 118
- error, RECON I/O 560

- ESAF (external subsystem attach facility) 13
- ESAF (External Subsystem Attach Facility)
  - accessing DB2 for z/OS 475
- establishing surveillance for XRF 606
- establishing XRF surveillance 606
- ESTAE process
  - bypassing affinity management
    - during 297
- ETO (Extended Terminal Option)
  - description 68
  - DFSLGNX0 (Logon exit routine) 327
  - DFSSGNX0 (Sign-on exit routine) 327
  - security 327
- event notification facility 186
- events
  - system-level tracing 380
- example
  - data sharing 245
  - display for structure full
    - threshold 183
  - explicitly stopping duplexing 188
  - RACF commands for authorizing CQS
    - registration 178
  - RACF commands to authorize
    - connection to CQS structures 178
  - RSR complex 687
- exclusive access 237
- exclusive intent
  - effect on scheduling 76
- exclusive transactions, Fast Path 73
- EXEC statement parameters
  - SSM 13
- execution procedures, tailoring
  - for data sharing 248
- exit
  - database open 497
- exit points
  - definition of 77
- exit routines
  - introduction to 77
  - RECON I/O
    - tracking changes to the
      - RECON 566
- exit routines for online change 403
- expedited message handling 73
- Extended Error Queue Element (EEQE)
  - in XRF complex 621
- Extended Recovery Facility (XRF) 617
  - active IMS
    - definition 586
    - failure 586
  - alternate IMS system 586
    - backup sessions 591
  - application programs 595
  - APPLID= keyword 647
  - ARM considerations 644
  - backup sessions 591
  - benefits 586
  - class-1 terminals 634
    - ownership of 642
  - class-2 terminals 635
  - class-3 terminals 636
  - comparison of MNPS and
    - USERVAR 590

- Extended Recovery Facility (XRF)
  - (continued)
  - complex
    - description of 588
    - initialization phase 608
    - synchronization phase 610
  - complex with one CPC 624
  - component roles 597
  - concepts 586
  - contributions
    - IMS 597
  - data set placement 656
  - DBCTL capabilities 591
  - defining IMS to VTAM 656
  - defining passwords to IMS 647
  - defining VTAM application name to
    - IMS 647
  - DFSHSBxx parameters 649
  - DFSMS contribution to XRF
    - process 601
  - DSE 586
  - DSE (dependent service element)
    - definition 597
  - example with one complex, two
    - CPCs 626
  - example with one complex, two CPCs,
    - and non-XRF IMS 626
  - example with two complexes, four
    - CPCs 629
  - example with two complexes, three
    - CPCs 628
  - forced takeover, VTAM 601
  - global resource serialization 644
  - hardware requirements 595
  - HSBID parameter 646
  - HSBMBR parameter 646
  - IMS master terminals 647
  - IMS secondary terminals 647
  - IMSplex 585
    - online change restriction without
      - RM 595
  - in an IMSplex 23
  - introduction 585
  - ISC link 633
  - JES considerations 644
  - licensed program requirements 595
  - limitations 632
  - LNK parameter for surveillance 649
  - local queue manager data sets 209
  - LOG parameter for surveillance 649
  - logging on 597
  - MNPS and USERVAR
    - comparison 590
  - MNPS parameter 646
  - MNPS= keyword 649
  - MNPSPW parameter 646
  - MNPSPW= keyword 649
  - MSC links 649
  - NCP
    - backup sessions 591
  - NCP contribution 605
  - NCP planning 643
  - network changes, initiating 616
  - NO parameter for surveillance 649
  - online change
    - IMSplex restriction without
      - RM 595

Extended Recovery Facility (XRF)  
(*continued*)

- operational requirements 595
- organization 624
- overview 585
- phases of processing 608, 624
- planning 632, 643
- RACF considerations 644
- RDS parameter for surveillance 649
- recommendation for DBRC 611
- recoverable service element (RSE) 586, 649
- requirements 588
  - hardware 595
  - software 595
- RSENAME= keyword 649
- RSR and 681
- RSR, comparison with 680
- signals from active IMS
  - IMS system log 597
  - ISC link 597
  - RDS 597
- software requirements 585
- specifying backup option 636
- specifying IMS.PROCLIB members 649
- SSP contribution 605
- surveillance
  - change 597
  - establishing 606
  - options 649
  - starting 597
  - stopping 597
- SWITCH= keyword for XRF 649
- sysplex
  - global online change 449
- system definition 645
- system definition macros for XRF 646
- takeover
  - causes 586
  - criteria for 606
  - definition 586
  - user perspective 638
- takeover conditions 649
- tasks performed on former active system
  - code maintenance 620
  - hardware configuration changes 620
  - performing hardware maintenance 620
  - performing library maintenance 620
  - preventive maintenance 620
  - testing backup IMS procedures 620
- terminals in 633
- terminology 586
- tracking phase
  - communication network status 611
  - database status 611
  - message queue status 611
  - MFS pool status 611
- USERVAR parameter 646

Extended Recovery Facility (XRF)  
(*continued*)

- USERVAR table, VTAM
  - defining IMS to VTAM 656
- USERVAR= keyword 649
- VTAM
  - backup sessions 591
  - forced takeover 601
- VTAM contribution to XRF process 601
- VTAM ownership affecting terminal switching 642
- XRF (Extended Recovery Facility)
  - network changes, initiating 616
  - z/OS contributions 600
- Extended Recovery Facility (XRF))
  - tracking phase
    - dependent region status 611
  - extent of recovery, determining for RSR 672
  - external subsystem attach facility (ESAF) 13
  - External Subsystem Attach Facility (ESAF)
    - accessing DB2 for z/OS 475
  - external subsystems
    - Data Capture exit routines 78
  - external systems
    - accessing 475
- EXVR online parameter
  - use in tuning 412

**F**

- facility class 175
- FACILITY class 178
- failure
  - IRLM
    - restart 270
  - logging on after affinity 295
  - system recovery 287
- Fast Path
  - application programs
    - message-driven 72
  - area-level data sharing 233
  - buffer pools 419
  - concepts and terminology 72
  - database tracker 669
  - databases
    - DEDB 72
  - DBCTL considerations 73
  - dependent regions
    - description 72
  - expedited message handling 73
  - IFP 72
  - ISC and 73
  - message handling 73
  - message queue list structures
    - defining 215
    - example 215
  - monitoring 384
  - MSC and 73
  - online change considerations 403
  - processing 72
  - registering databases and DEDB areas 499

Fast Path (*continued*)

- security considerations
  - DB/DC 321
- security considerations in DBCTL 344
- transactions
  - exclusive 73
  - recommendation 232
- Fast Path DEDB
  - record in the RECON data set 549
- Fast Path Log Analysis utility (DBFULTA0) 101, 349
- FDBR
  - as IMSplex component 23
- FDBR (fast database recovery)
  - DB/DC environment 10
  - DBCTL environment, in a 12
  - definition 10
  - FDRMBR parameter 10
- field-level sensitivity 339
- File Select and Formatting Print utility (DFSERA10) 349
- file select utility 177
- forced takeover
  - XRF 601
- formatting print utility 177
- forward database recovery
  - data sharing 509
  - non-data sharing 509
- forward recovery
  - data sharing 268, 509
  - definition 116
  - non-data sharing 509
- FREESPACE keyword
  - DEFINE CLUSTER keywords 542
- front-end IMS
  - in a shared-queues environment 205
- front-end switch planning 227
- functions of CQS 62
- fuzzy image copies
  - definition 516

**G**

- gap, definition 667
- GBL\_SERIAL\_PGM 173
- GDG (generation data group)
  - defining 36
  - model data set 37
- General IWAIT Time Events report
  - IMS Monitor (DB/DC) 751
  - IMS Monitor (DCCTL) 801
- Generalized Performance Analysis Reporting (GPAR)
  - use in detailed monitoring 375
- Generalized Trace Facility (GTF)
  - use in detailed monitoring 375
- generation data group (GDG)
  - defining 36
  - model data set 37
- generic resource
  - VTAM 295
- generic resource group
  - /START VGRS command 294



- generic resource group (*continued*)
  - /STOP VGRS command 294
  - dropping an IMS system 294
- generic resource groups
  - requirements 289
- generic resource name
  - specifying 291
    - GRSNAME= startup parameter 291
    - VGRS parameter 291
- generic resources
  - affinities
    - IMS managed 290
  - APPLID name 292
  - MNPS name
    - initiating a session 291
  - VTAM
    - planning for 289
    - restrictions on using 289
  - XRF
    - MNPS 292
    - USERVAR 292
- GENJCL command
  - generating JCL 493
- GENJCL.ARCHIVE command
  - log control requirements in the RECON data set 487
- GENJCL.IC
  - HALDB online reorganization considerations 531
- GENJCL.OIC
  - HALDB online reorganization considerations 531
- GENMAX parameter
  - image copy data sets for future use 521
  - resetting 569
  - specifying image copy requirements 521
- global callable services 300
- global command status 170
- global information
  - sysplex 173
- global online change
  - ACB library member 50, 454
  - cold start 463
  - command sequence 445
  - commands 445
    - INITIATE OLC 452
    - QUERY MEMBER 456
    - TERMINATE OLC 445
  - disabling
    - IMSplex 142
  - DRD
    - scenario 458
  - enabling
    - IMSplex 142
  - errors 457
  - functions 49
  - IMSplex 444
    - disabling 142
    - enabling 142
  - initializing 26
  - maintaining 142
  - mixed scope 143
  - MODBLKS=DYN
    - scenario 458

- global online change (*continued*)
  - MODBLKS=OLC
    - scenario 458
  - non-supported environments 440
  - overview 49
  - requirements 49
  - restrictions 142
    - without a resource structure 49
    - without an RM 49
  - RSR, and 142
  - scenarios 457
  - shutting down 142
  - starting 142
  - supported environments 437
  - TSO SPOC
    - output 458
    - utility 49
- Global Online Change utility 463
- Global Online Change utility (DFSUOLC0) 27, 449, 451
  - initializing OLCSTAT 49
- global resource information 169
- global resource serialization and XRF 644
- Global Resource Serialization macro 538
- global resources
  - types of
    - APPC descriptors 169
    - database names 169
    - DEDB area names 169
    - LTERMS 169
    - MSNAMEs 169
    - scheduled serial programs 169
    - transactions 169
    - user IDs 169
- global service group
  - record 552
- global service group (GSG) 669
- global status
  - availability during system failure 174
  - resetting to NULL 172
- GPAR (Generalized Performance Analysis Reporting)
  - use in detailed monitoring 375
- GPSB (generated program specification block) 18
- GRESTAE command 297
- groups
  - database 549
- GRS macro 538
- GRSNAME=
  - specifying a generic resource name 291
- GSAM (Generalized Sequential Access Method) 233
- GSG (global service group) 669
- GTF (Generalized Trace Facility)
  - use in detailed monitoring 375, 392
- GTF (generalized trace facility) trace 380
- guidelines
  - for application program preload 422
  - for placement of IMS system data sets 424

## H

- HALDB (High Availability Large Database)
  - DBRC commands supported 531
  - dynamic allocated data sets 32
  - naming conventions 83
  - online reorganization
    - DBRC considerations 531
  - REORG record 553
- HALDB (High Availability Large Databases)
  - master (DSPDBHRC)
    - DB header record 549
  - partition 549
  - partition DB record (DSPDBHRC) 549
  - partition DBDS records (DSPDSHRC) 549
  - partition record (DSPPTNRC) PHDAM 549
  - types of DBDSs 549
- hardware
  - XRF requirements 595
- hardware requirements
  - CQS 62
- header record
  - RECON data set 546
- high availability large database (HALDB)
  - online change 443, 447
- High Availability Large Database (HALDB)
  - dynamic allocated data sets 32
  - naming conventions 83
  - Partition Selection exit routines 78
  - REORG record 553
- High Availability Large Databases (HALDB)
  - master (DSPDBHRC)
    - DB header record 549
  - partition 549
  - partition DB record (DSPDBHRC) 549
  - partition DBDS records (DSPDSHRC) 549
  - partition record (DSPPTNRC) PHDAM 549
  - types of DBDSs 549
- hints and tips
  - using DBRC 567
    - locating the last SLDS Stop Time in the RECON data set 567
  - PRILOG compression not working 557
  - PRILOG record sizes 557
- hyperspace buffers, VSAM Buffer Pool //DFSSTAT 812
  - IMS Monitor (DB/DC) 753
  - IMS Monitor (DBCTL) 780
- HISAM Reorganization Unload utility (DFSURUL0)
  - for backup 523
  - introduction to 111
- HLQ parameter 120
- HOST macro 141
- HSBID parameter
  - used for XRF 646

HSBMBR parameter  
     used for XRF 646  
 HSSP data set 523  
     database registered with DBRC 523

## I

I/O prevention  
     by availability manager 614  
     by availability manager (AVM) 600  
     by operator 614, 618  
     definition 591  
     description 600  
     ensuring completion 600, 618  
     processing 600  
     XRF takeover 600  
 I/O resource contention  
     removing contention  
         tuning 428  
 I/O subsystem  
     configuration 425  
     dynamic monitoring 428  
     trade-offs  
         IMS and paging 420  
 I/O toleration  
     definition 615  
     operator view of 615  
 ICDSN parameter commands  
     creating for future use 520  
     duplicate, naming convention 515  
     maximum number of  
         generations 521  
     naming convention 514  
     record 553  
     recovery period 521  
     reusing 522  
 ICMD call  
     security overview 315  
 IDCAMS  
     commands  
         DEFINE 196  
         DELETE 196  
         REPRO 196  
 IFP (IMS Fast Path) 72  
     application program 8  
     message-driven program 8  
     non-message-driven program 8  
     region  
         characteristics 8  
         utility 8  
 ILDS (Indirect List Data Set)  
     as DBDS group 529  
     relation to DBRC 494  
 ILS (isolated log sender) 667  
 image copies  
     fuzzy 516  
     number of, specifying 521  
 image copies and data sharing 237  
 image copy  
     data sharing 260, 263  
     definition 111  
     with exclusive control 263  
 image copy data set  
     recovery period 521  
 image copy data sets  
     duplicate  
         naming conventions 515

image copy data sets (*continued*)  
     naming conventions 514  
 image copy group record 553  
 image copy purge times 104  
 importance  
     assigning to address space 372  
     business  
         Workload Manager 372  
     relative 372  
 IMS  
     dead-letter queue 229  
 IMSabend as cause of XRF  
     takeover 613  
 IMS Application Menu 41  
     overview 119  
     starting 119  
 IMS commands  
     /CHANGE SURVEILLANCE  
         XRF 597  
     /DISPLAY  
         for monitoring 375  
     /ERESTART BACKUP 608, 621  
     /START IMS 608  
     /START SURVEILLANCE  
         XRF 597  
     /STOP SURVEILLANCE  
         XRF 597  
     /SWITCH SYSTEM  
         planned takeover 613, 620  
         takeover process 613, 618  
     /SWITCH SYSTEM FORCE 613  
     /UNLOCK SYSTEM  
         ensuring database integrity 600  
         process during takeover 615  
 IMS components  
     BPE 33  
 IMS Connect  
     entry for z/OS PPT 128  
 IMS Control Center 45  
 IMS data set placement  
     requirement for XRF 595  
     XRF process, contribution to 597  
 IMS DataPropagator 78  
 IMS master terminals  
     defining for XRF 647  
 IMS Monitor 362, 378  
     DBCTL  
         timed events summary 768  
     description 388  
     distribution definition values 761  
     for detailed data 375  
     for performance analysis 407  
     multiple systems, using with 369  
     reports  
         Call Summary 362  
     reports for DCCTL  
         output sequence 788  
         overview 788  
         units of measure 788  
     testing 362  
     timed events summary  
         DBCTL 768  
 IMS Monitor Report Print utility  
     controlling 760  
 IMS Monitor reports  
     output  
         DB/DC 736

IMS Monitor reports (*continued*)  
     overview 733, 736  
     sequence of report output 736  
     units of measure in 736  
 IMS Monitor Reports  
     Buffer Pool Statistics 737  
     Call Summary 744  
     Communication Summary 754  
     Communication-Wait 755  
     Database Buffer Pool 753  
     DB/DC  
         Transaction Queueing 752  
 DBCTL  
     Call Summary 776  
     Deadlock Event Summary 781  
     Intent Failure Summary  
         Report 772  
     Latch Conflict Statistics 781  
     output selection options 770  
     overview 769  
     Pool Space Failure Summary 781  
     Program I/O 778  
     Program Summary 776  
     Programs by Region 772  
     Region Summary 772  
     Run Profile 770  
     sequence of report output 769  
     System Configuration 770  
     units of measure 769  
     verifying report occurrences 770  
     VSAM Buffer Pool 780  
 DCCTL  
     Call Summary 795  
     Communication Summary 803  
     Communication Wait 804  
     Distribution Appendix report 806  
     General IWAIT Time Events 801  
     internal resource use,  
         monitoring 804  
     Latch Conflict Statistics 804  
     Line Functions 803  
     Message Format Buffer Pool 800  
     Message Queue Pool 801  
     output selection options 789  
     overview 787  
     Pool Space Failure Summary 804  
     Program I/O 797  
     Program Summary 795  
     Programs by Region 791  
     Region and Jobname 789  
     Region Summary 791  
     Region Wait 791  
     reports that do not apply to 787  
     Run Profile 789  
     System Configuration 789  
     Transaction Queueing 802  
     verifying report occurrences 789  
 Deadlock Event Summary 755  
 Distribution Appendix report 783  
 Distribution-Appendix report 760  
 General IWAIT Time Events 751  
 Intent Failure 743  
 Latch Conflict Statistics 755  
 Line Functions 754  
 Message Format Buffer Pool 750  
 Message Queue Pool 750  
 MSC Queueing Summary 766

IMS Monitor Reports *(continued)*

- MSC Summaries 765
- MSC Traffic 764
- output selection options 739, 770
- Pool Space Failure Summary 755
- Print Program and MSC 763
- Program I/O 746
- Program Summary 744
- Programs by Region 740
- Region and Jobname 737
- Region Summary 740
- Region Wait 740, 772
- Run Profile 737
- System Configuration 737
- System Configuration report
  - adding to 770
- verifying report occurrences 739, 770
- VSAM Buffer Pool 753

IMS Monitor timed events

- checkpointing 733
- description 733
- DL/I call NOT-WAIT times
  - DB/DC 735
  - DCCTL 787
- during message input 733
- elapsed execution 733
- idle for intent 733
- NOT-WAIT time
  - DCCTL 787
- schedule of first DL/I call 733
- scheduling and termination 733
- summary
  - DB/DC 735
- trace intervals
  - DB/DC 737
  - DBCTL 770
  - DCCTL 789
- wait time 733
- wait-for-input (WFI)
  - DB/DC 733, 746
  - DCCTL 797

IMS procedure

- security 339

IMS region types

- batch message processing 8
- control region 8
- Fast Path 8
- message processing 8

IMS reports

- using 709

IMS secondary terminals

- defining for XRF 647

IMS shutdown 357

IMS SOAP Gateway

- using with IMS TM 475

IMS Spool API 465

- AFP 465
- application requirements 465
  - JES print data sets 465
- CHNG call
  - description 468
  - with OUTN option 468
  - with PRTO option 468
  - with TXTU option 468
- design considerations 465
- dynamic output 468
- JES as a Data Manager 468

IMS Spool API *(continued)*

- native terminal support 465
- OEM print server 465
- operational considerations 465
- OUTN option 468
- output data set 468
- print data set characteristics 468
- PRTO option 468
- SETO call 469
- TXTU option 468

IMS system

- administration
  - introduction 3
- introduction 3

IMS system access

- declaring
  - for batch systems 238
  - for online systems, changing online 238
  - for online systems, overview 238
  - online databases that share data 237

IMS system administration

- introduction 3

IMS system definition

- class-1 terminals, defining 648
- class-2 terminals, defining 648
- coding system definition macros for XRF 646
- defining master and secondary terminals for XRF 647
- macros with XRF-related keywords for terminals 648

IMS system log

- as surveillance mechanism 611
- during takeover 614
- during XRF takeover 597
- establishing as surveillance 606
- failure, XRF limitations 591
- override XRF takeover indication 597
- placement of for XRF 656
- recording
  - message queues 611
  - MFS pool loading 611
  - session information for XRF 597
  - status of dependent region activities 611
- tracking, used for 611
- tracking, XRF 597
- with USERVAR backup sessions 591

IMS system, documenting 81

IMS TM

- callout function 475
- IMS SOAP Gateway, using with 475

IMS TM Resource Adapter

- using with IMS TM 475

IMS Transaction Manager

- callout function 475

IMS.ACBLIB library

- allocating data sets dynamically 399

IMS.PGMLIB

- optimizing retrieval 412
- PSB names as members 70

IMS.PROCLIB

- DFSFIXxx member 611
- DFSHSBxx member
  - parameters 613

IMS.PROCLIB *(continued)*

- external subsystems 15
- OLDSDEF statement (PROCLIB) 94
- XRF parameters 649

IMSASAP II reports

- testing 362
- tuning 407

IMSCTF macro

- DFSCMC10 module 129

IMSGEN macro statement, security options 342

IMSplex

- address spaces participating 39
- ARM 705
- automatic RECON loss
  - notification 51
- batch environment 52
- coexistence with MSC 22
- command environment, type-2 68
- commands 39, 152
  - benefits 152
  - format 152
- compared with Parallel Sysplex 19
- components 19
  - optional 52
- components, required 52
- concurrent database access 233
- configuration recommendations 54
- configuring 52
- CQS 52
- CSL 38
  - configuration 52
  - configurations 54
  - recovery 705
- CSL (Common Service Layer)
  - monitoring 144
- CSLZQRY 144
  - DSECT function 144
  - STATS function 144
- data sharing 233
- DBRC support of 503
- defining 23
- definition 19, 39
- diagnosing problems 145
- displaying information 456
- DLISAS 52
- ensuring consistency 296
- FDBR 23
- generic resource group
  - consistency 296
- global online change 26, 444
  - ACB library member 50, 454
  - benefits 49
  - disabling 142
  - enabling 142
  - functions 49
- illustration 39
- IMS control region
  - defining 24
- IMS system services used 21
- maintaining availability for users 585
- Master IMS control region 49
- member
  - definition 19
- members
  - OM 39
  - RM 39

- IMSpIex (*continued*)
  - members (*continued*)
    - SCI 39
  - modifying resources 437
  - ODBM
    - configuring 24
    - initializing 24
  - OM
    - defining 25
    - OM API 152
    - RACF 158
    - security user exit routine 168
  - parameters
    - global status 173
  - PLEXPARM= parameter 170
  - recovery 705
  - resource structure 53
    - message destination resource 172
    - recommendation 44
    - recommendations 53
  - RM
    - defining 25
    - recommendations 53
  - RSR 23
  - SCI
    - initializing 25
  - security 30
    - example 175
    - type-2 commands 158
  - serial application processing 210
  - shared queues 52
    - using a resource structure 44
  - simplified 25
  - single system configuration 54
  - SPOC 152
    - commands 152
  - SPOC (single point of control) 45
  - tailoring 23
  - terminating online change 456
  - TM resources
    - managing 299
  - type-2 command environment 68
  - typical configuration 54
  - without RM (resource manager) 444, 451
  - XRF 23, 585
  - XRF (Extended Recovery Facility)
    - online change restriction without RM 595
  - z/OS system services used 21
- IMSpIex component 19
- IMSpIex member 19
- in-flight transactions
  - definition for XRF 597
  - processing during takeover 615
- inactive ACB library
  - resizing online 398
- Index/ILDS Rebuild utility (DFSPREC0)
  - relation to DBRC 494
- INDEXED keyword
  - DEFINE CLUSTER keywords 542
- INIT call
  - specifying data sensitivity 71
- INIT SELF command 139
- INIT.ADS command 507
- INIT.DB command
  - defining databases 507
- INIT.DBDS command
  - defining databases 507
  - REUSE keyword 522
  - specifying image copy requirements 521
- INIT.RECON command
  - initializing the RECON data set 536
  - RECON data sets 505
  - recovering RECONS 560
- initialization
  - IMSpIex, in an 304
- Initialization exit routine (DFSINTX0)
  - using with ETO 327
- initialization phase
  - description 608
  - illustration of 608
  - procedures for operators 608
- initializing
  - DBRC
    - IMSCTRL Macro 505
    - RECON data set 505
  - INITIATE OLC 49
  - INITIATE OLC command 445, 452
    - error handling 454
    - return and reason codes 453
  - initiating a takeover
    - procedure for 613
    - process during takeover 613
    - purpose of 620
  - initiating network changes at takeover 616
- input/output error processing 560
- installation
  - considerations for IMS 125
  - problems, preventing 125
  - service
    - attention notice 195, 200
    - preventive 195
  - tasks
    - data sharing 246
  - z/OS interface 125
- installation issues 201
- instance, as part of RSR name 669
- Intent Failure Summary Report
  - IMS Monitor (DB/DC) 743
  - IMS Monitor (DBCTL) 772
- interest
  - registering interest
    - concept 206
- interface considerations, attention
  - notice 125
- interface modules 129
- internal resource lock manager (IRLM)
  - entry for z/OS PPT 129
  - z/OS PPT entry 129
- Internal Resource Lock Manager (IRLM)
  - data sharing 499
  - database integrity 499
  - in data sharing 239
- interpret table
  - example of 601
  - initialization of 601
  - XRF's use of 601
- interpreting
  - //DFSSTAT reports 811
  - IMS Monitor Reports
    - DBCTL 767
- interpreting (*continued*)
  - IMS Monitor Reports (*continued*)
    - DCCTL 787
  - Log Transaction Analysis Reports 825
  - Intersystem Communication 73
  - interval value
    - setting 606
  - interval value for XRF surveillance 606
  - introduction 79
  - IRLM
    - activate
      - execution JCL 238
      - IRLMNM parameter 238
    - CSA and ECSA, use of 241
    - database-level sharing 238
    - failure 286
      - restart 270
    - MAXCSA= keyword 241
    - PC= keyword 241
    - system initialization 242
    - z/OS subsystem 242
  - IRLM (internal resource lock manager)
    - as part of an IMSpIex 21
    - entry for z/OS PPT 129
    - execution parameters 240, 241
    - installation of 240
    - monitoring 252, 352
    - recovery 269
    - sharing modes 241
    - start 249
    - stopping 250
    - tracing 380
    - z/OS PPT entry 129
  - IRLM (Internal Resource Lock Manager)
    - activate
      - for batch systems 247
    - as cause of takeover 613
    - data sharing 499
    - data sharing, role in 239
    - database integrity 499
    - failure 613
    - failure as cause of XRF takeover 591
    - illustration of during XRF takeover 618
    - IRLM IVP subset
      - APF authorization 137
      - dump formatting module 137
      - naming suggestions for VTAM interface 137
      - PPT entry requirements 137
      - subsystem names 137
    - processing during XRF takeover 618
    - used as local lock manager in XRF 618
    - with batch backout 509
  - IRLM structure
    - calculating size 280
  - ISC (Intersystem Communication)
    - Fast Path and 73
    - link between active and alternate IMS system
      - as surveillance mechanism 611
      - definition 597
      - establishing as surveillance 606
      - illustration of 633
      - initializing 608

ISC (Intersystem Communication)  
     *(continued)*  
     link between active and alternate IMS system *(continued)*  
         recommendation for XRF 633  
         takeover condition 597  
         used for synchronization 610  
 isolated log sender (ILS) 667  
 IVP  
     service 202

## J

JCL (job control language)  
     generated by DBRC 480  
     skeletal execution members  
         definition 480  
         tailoring for utilities 487  
 JCL allocation  
     RECON data set 541  
 JES (job entry subsystem)  
     failure, XRF limitations 591  
     planning considerations 644  
 JES configuration 126  
 JMP (Java message processing)  
     application security 321  
     security for applications 321  
 job control language (JCL)  
     generated by DBRC 480  
     skeletal execution members  
         definition 480  
 job entry subsystem  
     failure, XRF limitations 591  
 job summary report 380  
 JOBJCL  
     skeletal JCL execution member 493

## K

KBLA (Knowledge-Based Log Analysis)  
     defining selection of logs 841  
     external log processing utility  
         accessing 844  
         invoking 839  
         ISPF interface 838  
         maintaining  
             option 0, using 841  
         running jobs against IMS log records  
             option 1 842  
             option 2 842  
             option 3 842  
             option 4 842  
         utilities overview 833  
 keyboard shortcuts xiv  
 KEYS keyword  
     DEFINE CLUSTER keywords 542  
 keywords  
     MAXAPPL 139  
 Knowledge-Based Basic Formatting Print  
     routine (DFSBLA3) 349  
     overview of 833  
 Knowledge-Based Formatting Print  
     routine (DFSBLAK) 349  
     overview of 833  
 Knowledge-Based Log Analysis (KBLA)  
     defining selection of logs 841

Knowledge-Based Log Analysis (KBLA)  
     *(continued)*  
     external log processing utility  
         accessing 844  
         invoking 839  
         ISPF interface 838  
         maintaining  
             option 0, using 841  
         running jobs against IMS log records  
             option 1 842  
             option 2 842  
             option 3 842  
             option 4 842  
         utilities overview 833  
 Knowledge-Based Log Analysis (KBLA)  
     utilities 349  
 Knowledge-Based Summary Formatting  
     Print routine (DFSBLAS) 349  
     overview of 833

## L

Latch Conflict Statistics report  
     IMS Monitor (DB/DC) 755  
     IMS Monitor (DBCTL) 781  
     IMS Monitor (DCCTL) 804  
 legal notices  
     notices 847  
     trademarks 847  
 LGEN subparameter  
     preprocessor 402  
     stage 1 processing 402  
 libraries  
     active and inactive 438  
 Library Lookaside  
     program libraries 422  
 Line and Terminal report  
     example 826  
 line failure  
     XRF limitations 591  
 Line Functions report  
     IMS Monitor (DCCTL) 803  
 Line-Functions report  
     IMS Monitor (DB/DC) 754  
 link queuing time assessments 766  
 linking IMS to z/OS 129  
 LNK parameter for surveillance 649  
 LNK parameter in member  
     DFSHSBxx 606  
 local online change  
     cold start 463  
     command sequence 441  
     supported environments 437, 440  
 local online change function  
     overview of 438  
 local processing  
     defined 205  
 Local Shared Resources (LSR)  
     option 542  
 lock structure  
     calculating size 280  
 locked messages  
     on a shared queue 208  
 locks for uncommitted database  
     changes 611  
 log  
     archiving 99

log *(continued)*  
     change accumulation 103  
     compression 110  
     concepts 74  
     content of 110  
     controlling the characteristics of 353  
     Data Capture exit routine 110  
     data set  
         specifying choices 106  
     data sets 94  
     data sharing 265  
     failure, XRF limitations 591  
     introduction to 92  
     OLDS (online log data set) 94  
     records  
         printing 349  
         reports 349  
     reducing 110  
     RLDS 97  
     SLDS 96  
     system utilities 349  
     tracing 101  
     use in restart 74  
     WADS 95  
     z/OS 98  
 log allocation record 554  
 Log Archive utility (DFSUARC0) 95, 96  
     archiving  
         manual 99  
     customized with exit routines 101  
     description 487  
     OLDS archive 99  
 log data sets  
     moved, notifying DBRC of 574  
 log data sets, records 546  
 log information 485  
 LOG keyword  
     DEFINE CLUSTER keywords 542  
 log management, RSR 685  
 LOG parameter  
     for surveillance 649  
     in member DFSHSBxx 606  
 log records  
     analyzing 825  
     deleting from RECON 558  
     OM audit log  
         record format 155  
     OM, printing 155  
 log router  
     use with RSR 666  
 log stream  
     z/OS  
         defining 217  
 Log Transaction Analysis utility  
     (DFSILTA0) 349  
     as tuning aid 407  
     multiple systems, using with 369  
     reports produced  
         description 828  
         Log Analysis report 828  
     save queue times 416  
     used for monitoring 375  
 log volumes  
     specifying for change  
         accumulation 526



- log-related commands
  - API requests
    - Log query 488
    - OLDS query 488
  - CHANGE.PRILOG 488
  - CHANGE.RECON 488
  - CHANGE.SECLOG 488
  - DELETE.LOG 488
  - GENJCL.ARCHIVE 488
  - GENJCL.CLOSE 488
  - LIST.LOG 488
  - NOTIFY.PRILOG 488
  - NOTIFY.SECLOG 488
- LOGCHAR macro 601
- logging 67
  - accumulating logs using DFSUCUM0 525
  - condensing logs using DFSUCUM0 525
  - dual 106
    - definition 94
  - IMS events 485
  - overview 92, 485
  - reducing, for Fast Path 101
  - single 106
  - states 106
- logging in MSC 383
- logging on to alternate IMS system 601
- logging on to alternate subsystem 597
- logging on to XRF IMS
  - illustration of 601
  - processing of logon message 601
- logic review 82
- logical terminal (LTERM)
  - as a TM resource in a sysplex 301
  - definition to RM 301
  - name uniqueness 299, 301
  - resource type consistency 300
  - shared queues
    - when IMS registers interest 207
- logical unit definitions for VTAM 139
- LOGMODE parameter 139
- LOGON APPLID command 601
- Logon exit routine (DFS LGNX0)
  - using with ETO 327
- logon message
  - example of 601
  - illustration of 601
  - processing of 601
- logon procedure
  - generic resource name 292
  - IMS APPLID name 292
- logs
  - control records in RECON
    - changing 486
- LSO (local storage option)
  - for IRLM control blocks 241
- LSR (Local Shared Resources)
  - option 542
- LTERM (logical terminal)
  - as a TM resource in a sysplex 301
  - definition to RM 301
  - name uniqueness 299, 301
  - resource type consistency 300
  - shared queues
    - when IMS registers interest 207

- LU 6.2
  - command authorization 313
- M**
- macro keywords
  - RACFCOM on SECURITY macro 331
  - RACFTERM on SECURITY
    - macro 331
  - RCLASS on SECURITY macro 331
  - SECCNT on SECURITY macro 331
  - SECLVL on SECURITY macro 315, 331
  - TERMNL on SECURITY macro 315, 331
  - TYPE on SECURITY macro 331
- macros
  - HOST 141
- macros for XRF system definition 646
- MADS (multiple area data set)
  - and RSR 692
- main storage database 72
- main storage database (MSDB)
  - status as global resource 170
- maintaining global resource
  - information 169
- maintaining separate copies of data sets
  - XRF 656
- maintenance
  - APPLY status 201
  - assessing readiness before applying 193
  - considerations 191
  - preventing regression 201
  - production system 194
  - recommendations 193
  - sample process 194
  - test cycle 194
  - verifying with test system 359
- maintenance issues 201
- managing structure usage 182
- managing system definition
  - online change 402
- marooned data, definition 667
- master database
  - definition 665
- Master IMS control region 49
- master terminal
  - defining
    - for VTAM 647
  - definition 9
  - IMSplex, in an 152
  - security
    - enforcing security options 315
    - use of signon verification 315
- MAXAPPL keyword 139
- MAXCSA= parameter 241
- merging logs
  - data sharing 526
- message
  - CQS0033A 177
  - CQS0034A 187
  - CQS0205E 182
  - expedited message handling 73
  - Fast Path exclusive 73

- message (*continued*)
  - flow
    - in a shared-queues environment 208
  - format
    - options 412
  - input 6
  - IXC585E 183
  - IXC586I 183
  - locked on shared queue
    - concept 208
  - MTO 210
  - output 6
  - queue
    - built-in alternate 610
    - processing during takeover 614
    - recorded on log 611
    - transactions and 6
  - message format buffer pool
    - optimizing 417
  - Message Format Buffer Pool Report
    - IMS Monitor (DB/DC) 750
    - IMS Monitor (DCCTL) 800
  - Message Format Service (MFS)
    - MFSTEST mode 364
    - MFSTEST=YES on IMSGEN
      - macro 365
    - test formats 365
    - testing formats online
      - MFSTEST mode 364
      - online execution
        - requirements 365
      - system definition
        - requirements 365
  - message processing program (MPP)
    - scheduling 70
  - message queue
    - backup 112
    - deleting 219
  - message queue list structures
    - defining
      - Fast Path 215
    - defining in a shared-queues environment 215
  - message queue pool
    - optimizing 416
  - Message Queue Pool Report
    - IMS Monitor
      - DB/DC 750
    - IMS Monitor (DCCTL) 801
  - message routing
    - performance
      - analysis 383
  - message switching
    - shared queues
      - Destination Creation exit routine (DFSINSX0) 209
      - undefined destination
        - Destination Creation exit routine (DFSINSX0) 209
  - message-driven program
    - IFP application programs 8
    - MPP 8
  - message-driven programs 72
  - messages
    - Queued But Not Sent report
      - example 827

- Messages
  - Program To Program Report
    - example 827
    - report example 828
- messages and codes
  - Base Primitive Environment (BPE) 33
- MFS (Message Format Service)
  - logical paging and XRF takeover 638
  - MFSTEST mode 364
  - MFSTEST=YES on IMSGEN
    - macro 365
  - pool loading recorded on log 611
  - test formats 365
  - testing formats online
    - MFSTEST mode 364
    - online execution
      - requirements 365
    - system definition
      - requirements 365
- MFS formats
  - changing 443, 446
- MFS Service utility
  - procedure MFSRVC for index 430
  - use in tuning 416
- MFSTEST mode 364
- migrating
  - tasks 204
- milestone, definition 668
- MNPS name
  - in a generic resources
    - environment 291
- MNPS parameter
  - used for XRF 646
- MNPS= keyword for XRF 649
- MNPSPW parameter
  - used for XRF 646
- MNPSPW= keyword for XRF 649
- MODBLKS
  - system definition
    - needed for online change
      - function 438
- MODE parameter 139
- MODETBL parameter 139
- MODIFY (/MODIFY) command
  - PREPARE 450
  - using for online change 441
- MODIFY command
  - VTAM 606
- modifying
  - system design 395
- modifying resources online 437
- MODSTAT
  - DD statement 49
- MODSTAT data set
  - placement of in XRF 656
- MODSTAT2 49
- monitor trace interval
  - IMS Monitor Report
    - DBCTL 770
    - DCCTL 789
  - IMS Monitor Reports
    - DB/DC 737
- Monitor-Overhead report
  - description 730
  - fields in the report 730
- Monitor, IMS 362
- monitoring
  - /DISPLAY command 375, 392
  - application program elapsed time
    - DCCTL 795
  - data-sharing systems 252
  - database buffers
    - DB/DC 753
    - DBCTL 780
  - DBCTL 392
  - DBCTL considerations 389
  - dependent regions
    - DB/DC 740
    - DBCTL 772
    - DCCTL 791
  - DISPLAY (/DISPLAY) command 375, 392
  - dynamic 428, 430, 431
  - establishing objectives 391
  - Fast Path systems 384
  - frequency 375
  - I/O for application program DL/I
    - calls
      - DCCTL 797
  - IMS 349
  - IMS Monitor 378
  - in test environments 362
  - internal resource usage
    - DB/DC 755
    - DBCTL 781
    - DCCTL 804
  - IRLM activity 252, 352
  - line activity
    - DB/DC 754
    - DCCTL 803
  - message handling
    - DB/DC 755
    - DCCTL 804
  - message queue handling
    - DB/DC 746, 750
    - DBCTL 778
    - DCCTL 801
  - MFS activity
    - DB/DC 750
    - DCCTL 800
  - multiple systems 369
  - performance objectives
    - Workload Manager 371
  - procedures 369, 389
  - reasons for 369
  - strategies 369, 389
  - structure 254, 352
  - system 352
  - tools 378
  - using frequency distribution
    - DB/DC 758
    - DBCTL 783
    - DCCTL 806
  - your system 369
- MPP (message processing program)
  - region
    - characteristics 8
    - starting 8
- MSC
  - security
    - transactions 324
    - when RACF and exit routines are
      - called 324
- MSC (multiple systems coupling)
  - defining with the CTC option 132
  - remote processing
    - defined 205
  - shared queues
    - when IMS registers interest 208
  - shared-queues environment, in a 230
- MSC (Multiple Systems Coupling)
  - determining cross-system
    - queuing 764
  - Fast Path and 73
  - IMS Monitor Report Print
    - Program 763
  - Interpreting
    - IMS Monitor MSC Reports 763
  - interpreting Distribution Appendix
    - output 809
  - links
    - defining for XRF 649
  - monitoring 369
  - MSC Queuing Summary report 766
  - MSC-Queuing Report
    - interpreting 763
  - MSC-Summaries Report
    - assessing queue sizes 765
    - content 765
    - example 765
    - interpreting 763
  - MSC-Traffic Report
    - content 764
    - cross-system queuing,
      - determining 764
    - example 764
    - performance information 383
  - MSC Queuing Summary report
    - assessing link queuing times 766
    - example 766
- MSC security
  - Security Reverification exit routine
    - (DFSCTSE0) 326
  - Signon/off Security exit routine
    - (DFSCSGN0) 326
  - specifying user IDs 325
  - TM and MSC Message Routing and
    - Control exit routine
      - (DFSMSCE0) 325
  - Transaction Authorization exit routine
    - (DFSCTRN0) 326
- MSC Summaries Report
  - IMS Monitor (DB/DC) 763
- MSCSEC=
  - MSC transaction security 324
- MSDB (main storage database)
  - definition of 610
  - Fast Path considerations 72
  - loaded by alternate IMS system 610
  - not supported for data sharing 233
  - placement of in XRF 656
  - status as global resource 170
- MSGQ primary list structure 215
- MSNAME
  - as a TM resource in a sysplex 301
  - RM definition 301
- MSPLINK macro
  - BACKUP keyword 649
- MTO (master terminal operator)
  - description 355

- MTO (master terminal operator)
  - (continued)
  - messages 210
- MTO (Master Terminal Operator)
  - backup 9
  - responsibilities 9
  - secondary 9
- multiple area data set (MADS)
  - and RSR 692
- multiple cold starts
  - test environment 575
- multiple system failure
  - synchronizing IMS restart 287
- Multiple Systems Coupling 73
- Multiple Systems Coupling (MSC)
  - coexistence with an IMSplex 22
  - interpreting Distribution Appendix
  - output 809
  - links
    - defining for XRF 649

## N

- NAME keyword
  - DEFINE CLUSTER keywords 542
- name type 300
- name uniqueness
  - resource
    - disabling enforcement 299
- naming conventions
  - change accumulation data sets 515
  - data sharing, and 236
  - DBRC data sets 513
  - duplicate image copy data sets 515
  - establishing 83
  - examples of 83
  - HALDB 83
  - image copy data sets 514
  - RSR, for 669
  - SSIDs in RECON SSYS records,
    - for 572
  - SSIDs processed by batch backout,
    - for 572
  - suggestions 83
- NCP (Network Control Program)
  - communications network tuning 423
  - failure, XRF limitations 591
  - requirement for XRF 588, 595
  - storage considerations 643
  - switching sessions at takeover 591, 616
  - XRF process, contribution to 605
- NCP considerations 141
- NCP delay 141
- network
  - changes to definition 400
  - documentation 85
  - ensuring readiness 362
- network changes during XRF
  - takeover 616
- network operators at XRF complex
  - adding entries in USERVAR
    - table 601
  - MODIFY USERVAR command 601, 616
  - responsibilities resulting from
    - XRF 595

- network operators at XRF complex
  - (continued)
  - updating entries in USERVAR
    - table 606
- NO parameter for surveillance 649
- node name
  - uniqueness 302
- non-message-driven program
  - BMP 8
  - IFP utility 8
- non-swappable work on alternate IMS
  - system 632
- nonstandard image copy data sets
  - description 524
- nonstandard macros 126
- NOREUSE parameter
  - CA data set 528
  - image copy data sets for future
    - use 520
- normal operations
  - alternate IMS system, tracking 611
  - alternate tracking during 597
- normal restart
  - with data sharing 512
- nothing-to-process message
  - change accumulation 573
- NOTIFY.IC command
  - HALDB considerations 531
- NOTIFY.UIC command
  - HALDB considerations 531
- notifying DBRC that log data sets have
  - moved 574
- NOWRITECHECK keyword
  - DEFINE CLUSTER keywords 542

## O

- OBTAIN macro 538
- occupancy, region
  - measuring 743
- ODBA
  - configuring ODBA application servers
    - for ODBM 149
- ODBA application programs
  - security considerations in 322, 344
- ODBM
  - message routing 149
- ODBM (Open Database Manager)
  - configuration requirements 54
  - configuring 24
  - initializing 24
- ODBA
  - configuring ODBA application
    - servers for ODBM 149
  - overview 41
  - RRS 148
  - security
    - overview 150
  - supported interfaces 41
- offline dump formatting module
  - DFSOFMD0 136
- OLC= parameter 26
- OLCSTAT
  - cold start 463
- OLCSTAT data set 26, 173
  - attributes 27
  - format 27
- OLCSTAT data set (continued)
  - header sample 27
  - recommended attributes 27
  - XRF (Extended Recovery Facility)
    - IMSplex restriction without
      - RM 595
- OLDS
  - block size 421
  - buffers 421
  - minimize system log I/O 421
- OLDS (Online Data Set)
  - closing in a test environment 575
- OLDS (online log data set)
  - access method 94
  - allocating 94
  - archive 99
  - archiving 94
  - archiving DBRC 487
  - blocksize 107
  - buffers 109
  - changing the characteristics of 109
  - closing from WADS 95
  - defining 94, 106
  - definition 92
  - degraded mode logging 106
  - devices 108
  - dual or single 106
  - environment 94
  - number of 107
  - RECON 94
  - records 546
  - requirement for XRF 588
  - reusing 99
  - size 107
  - stopping 94
- OLDSDEF statement 94
- OM 39
  - audit log
    - record format 155
  - command security 157
  - log records, printing 155
- OM (Operations Manager)
  - audit trail 153, 154
  - configuration requirements 54
  - defining 25
  - OM API 152
  - XRF 585
- online change
  - ACB member
    - associating 448
    - associating with DRD
      - commands 448
  - cold start 463
  - command sequence 441
  - controlling 259
  - DEDB 443, 447
  - description 437
  - for security changes 337
  - global 49
    - ACB library member 50, 454
    - after DBCTL standby emergency
      - restart 449
    - after XRF takeover 449
    - cold start 463
    - command sequence 445
    - errors 457
    - IMSplex 444



- online change (*continued*)
  - global (*continued*)
    - scenarios 457
    - terminating 456
  - global status 170
  - HALDB (high availability large database) 443, 447
  - initiating 452
  - local
    - cold start 463
    - command sequence 441
    - sysplex 450
  - performance considerations 444, 450
  - prepare command failure
    - scenario 461
    - timeout error 461
  - programs 442, 446
  - security changes for DBCTL 347
  - system definition 441, 445
  - system definition changes
    - allowed 402
  - transactions 442, 446
  - XRF 443, 448, 449
  - XRF (Extended Recovery Facility)
    - IMSplex restriction without RM 595
    - XRF complex, in a 656
- online change commands
  - INITIATE OLC 445
  - QUERY 445
  - TERMINATE OLC 445
- online change data sets 438
- online change function
  - overview of 438
- Online Change utility (DFSUOCU0) 27
  - maintaining copies of IMS data sets in XRF 656
- online change, local
  - commands 441
- Online Data Set (OLDS)
  - closing in a test environment 575
- Online Database Image Copy utility (DFSUICP0) 519
  - creating data sets for future use 520
  - description 516
  - execution recorded by DBRC 515
- online execution requirements for MFSTEST 365
- online forward recovery, definition 667
- online log data set (OLDS)
  - access method 94
  - allocating 94
  - archiving 94
  - blocksize 107
  - buffers 109
  - changing the characteristics of 109
  - closing from WADS 95
  - defining 94, 106
  - definition 92
  - degraded mode logging 106
  - devices 108
  - dual or single 106
  - environment 94
  - number of 107
  - RECON 94
  - records 546
  - reusing 99
- online log data set (OLDS) (*continued*)
  - size 107
  - stopping 94
- online reorganization
  - DBRC commands supported 531
  - maintaining recovery records 533
  - process description 533
  - recommendations 533
  - records used 531
  - REORG records 531
  - restrictions 533
  - RSR tracking system
    - in RECON data set 533
- online system
  - definition 5
- online testing
  - during system test 363
  - MFS formats 364
  - MFSTEST mode 364
  - with Batch Terminal Simulator 364
- online updater
  - reinstating 261
- Open Database Manager (ODBM)
  - administration 147
  - configuration requirements 54
  - configuring 24
  - initializing 24
  - ODBA 148
    - configuring ODBA application servers for ODBM 149
  - overview 41
  - registering clients 148
  - RRS 148
  - security
    - overview 150
  - supported interfaces 41
- operating IMS
  - choosing monitoring tools 378
  - controlling data sharing 352
  - IMS Monitor 378
  - monitoring IMS 349
  - processing system logs 349
  - recovering
    - connecting subsystems 353
    - disconnecting subsystems 353
    - modifying log data sets 353
    - tasks 350
    - tuning log data sets 353
  - tasks, list of 349
  - tracing
    - IMS trace facility 382
    - program isolation and lock 381
    - z/OS component trace service 380
    - z/OS generalized trace facility 380
- operational procedures, testing 361
- operations
  - automated 88
  - introduction 87
- Operations Manager (OM)
  - administration tasks 151
  - command routing 151
  - command security 157
  - configuration requirements 54
  - functions provided 44
  - overview 44
- operations task
  - overview of 87
- operators at XRF complex
  - bringing up
    - new alternate IMS system 621
    - third system as alternate IMS system 621
  - communicating with each other 614, 618
  - ensuring I/O prevention is complete 618
  - initializing IRLM 619
  - initiating a takeover
    - practical uses 620
    - procedure for 613
    - takeover process 613
  - manual control at takeover 613
  - performing I/O prevention manually 600, 618
  - reestablishing service on class-3 terminals 588
  - reinstating RACF reverify capability 644
  - replying GO to message AVM005A
    - alternate IMS system, when I/O prevention complete 614
    - in-flight transactions 615
  - replying GO to message AVM006E 600
  - requesting SNAPQ checkpoint 610
  - resetting the CPC 618
  - responsibilities resulting from XRF 595
  - starting dependent regions on alternate IMS system 610
  - takeover tasks 638
- OPNDST (/OPNDST) command 139
- optimizing
  - application program load 422
  - dispatching priority 410
  - I/O contention 434
  - IMS system data sets 424, 430
  - message format buffer pool 416
  - message format libraries 430
  - message queues 430
  - number of message regions 414
  - processing priority for IMS regions 428
  - PSB and DMB pools 416
  - scheduling/termination 375
  - with page fixing 412
- OSAM (Overflow Sequential Access Method)
  - buffer pools 419
  - cache structures
    - changing size of 283
  - structure
    - calculating size 282
- OSAM Sequential Buffering 416, 419
- OSAM-Buffer-Pool Report 813
- OTMA (open transaction manager access)
  - shared-queues environment 228
  - asynchronous messages 228
  - synchronous messages 228
- output from IMS Monitor reports
  - DB/DC 736

- output sequence and information from
  - IMS Monitor Report
    - DCCTL 788
- output sequence of reports
  - IMS Monitor Report
    - DBCTL 769
- output, printing with IMS Spool
  - API 465
- overflow
  - mode 183
  - processing 183
  - threshold 183
- overflow processing
  - structure 66
- Overflow Sequential Access Method (OSAM)
  - cache structures
    - changing size of 283
  - structure
    - calculating size 282

## P

- page fixing
  - DREF storage 412
  - expanded storage 412
  - for control region 412
  - using member DFSFIXxx 412
- parallel access to RECON
  - planning for 537
- parallel RECON access 143
- parallel session support 139
- Parallel Sysplex
  - compared with IMSplex 19
- parameters
  - ALTRESL 120
  - DBRC for online IMS 505
  - HLQ 120
  - IMS control region execution
    - parameters
      - defining 218
- partition DB record (DSPDBHRC)
  - HALDB
    - TYPE=PART 549
- partition DBDS records (DSPDSHRC)
  - HALDB
    - types of DBDSs 549
- partition record (DSPPTNRC)
  - HALDB
    - PHDAM 549
- partitioned configuration
  - shared queues 212
- partitioning
  - a shared-queues configuration 212
- PassTickets
  - security 311
- password protecting
  - /LOCK command 314
  - /SET command 314
  - /UNLOCK command 314
- password verification
  - disable 311
  - enable 311
- passwords 307
  - in ACB 647
  - masking 338

- performance
  - analysis tools
    - RMF II 407
  - choosing IMS options for 412
  - criteria 370
  - design variables 409
  - factors
    - for system initialization 410
  - objectives 370
  - goals 408
  - management
    - change 409
  - NCP (Network Control Program)
    - considerations 423
  - network considerations 423
  - objectives 370
    - Workload Manager 371
  - of MSC 383
  - overview 408
  - planning
    - shared-queues environment, in
      - a 426
  - reporting aids
    - as part of monitoring 375, 378
  - tuning aids
    - list of 407
  - VTAM considerations 423
- performance goals
  - Workload Manager 372
- performance management
  - change 409
- performance study
  - online system design 402
- phases of the XRF process 608
- physical security 338
- PI (program isolation)
  - in DB/DC environment 76
- PI keyword 381
- PL/I subroutines, preloading 422
- planned takeover
  - changing IMS system definition 620
  - definition of 613
  - example of 621
  - limitations of 620
  - operator procedures for 620
  - practical uses of 620
  - uses of 591
- planning
  - application change, example of 395
  - introduction 87
  - system definition changes 399
- PLEXARM parameter
  - GSTSAREA 173
  - GSTSDB 173
  - GSTSTRAN 173
- Pool Space Failure Summary report
  - IMS Monitor (DBCTL) 781
  - IMS Monitor (DCCTL) 804
- pool space failures
  - resolving 418
- Pool-Space-Failure-Summary report
  - IMS Monitor (DB/DC) 755
- post-takeover phase
  - description of 621
  - illustration of 621
- potential transactions, Fast Path 73

- PPT (Program Properties Table) for z/OS
  - entry for BPE 127
  - entry for CQS 127
  - entry for IMS 127
  - entry for IMS Connect 128
  - entry for IRLM 129
  - updating 126
- PPT entry requirements 137
- practical uses of the planned takeover 620
- prediction
  - explicit 410
  - implicit 410
- PREPARE command 442, 446
- preventative service
  - installing 195
  - IMSplex 200
- preventing structure full 182
- preventive service
  - Authorized Program Analysis Report (APAR) 191
  - Program Temporary Fix (PTF) 191
- PRILOG
  - compression, automatic 557
  - compression, diagnosing problems
    - with 557
  - compression, manual 557
  - record sizes 557
- PRILOG record
  - closing an open online 570
- Print Dump Exit Control Table 137
- printed output 465
- printers
  - associated 226
- priority of session recovery
  - class-1 terminals, defining 648
  - class-2 terminals, defining 648
  - default values for IMS support 646, 648
- priority of session recovery for XRF
  - defining MSC links 649
- private queue types managed by CQS 65
- procedures
  - including DBRC 505
- processing continuity 74
- processing intent 237
- processing, Fast Path 72
- production configuration
  - documentation 85
- production system
  - maintenance 194
- productivity aids
  - Batch Terminal Simulator for
    - testing 364
  - data dictionary 84
  - for monitoring 375, 378
- productivity tools
  - database 350
- program deadlock
  - how resolved 76
- Program I/O report
  - IMS Monitor
    - DCCTL 797
    - IMS Monitor (DBCTL) 778
- program isolation (PI)
  - in DB/DC environment 76

- Program Isolation Trace Report utility (DFSIRP0) 381
- program library, optimizing 433
- program load
  - preload option 422
- Program Properties Table (PPT) for z/OS
  - BPE entry 127
  - CQS entry 127
  - IMS Connect entry 128
  - IMS entry 127
  - IRLM entry 129
  - updating 126
- program scheduling
  - introduction 70
- Program Summary Report
  - IMS Monitor
    - DB/DC 744
    - DBCTL 776
    - DCCTL 795
- program temporary fix (PTF)
  - applying in XRF complex 620
- program testing using
  - SYSIN/SYSOUT 366
- Program-I/O report
  - description 722
  - distributed event default ranges 729
  - events, distributed 729
  - fields in the report 722
  - using the report 722
- Program-I/O Report
  - IMS Monitor
    - DB/DC 746
- program-to-program
  - shared queues
    - Destination Creation exit routine (DFSINSX0) 209
    - undefined destination
      - Destination Creation exit routine (DFSINSX0) 209
- Programmed Cryptographic Facility 341
- programs
  - GBL\_SERIAL\_PGM 173
  - serialized
    - GBL\_SERIAL\_PGM 173
    - sysplex, managing in 173
- Programs by Region Report
  - IMS Monitor
    - DB/DC 740
    - DBCTL 772
    - DCCTL 791
- protecting resources with RACF, DLISAS
  - procedure 339
- PSB (program specification block)
  - program name convention 70
  - requests
    - in CCTL regions 414
  - resource security 341
- PSB buffer pool
  - analyzing requirements 418
- PSB work buffer pool
  - storage requirements 418
- PST-Accounting Report 811
- PTF (program temporary fix)
  - using SMP/E 191
- PTFs
  - Recommended Service Upgrade (RSU) 192

- publications
  - changes xvii
  - organization xvii
  - titles xvii
- purge time
  - Database Change Accumulation 104
  - definition 526

## Q

- QASTSPE variable 137
- QCF (Queue Control Facility)
  - cold queues
    - requeuing 222
  - structures
    - monitoring 222
    - requeuing 222
    - testing IMS 367
- QUERY command 445
- QUERY MEMBER command 456
- QUERY OLC command 456
- queue
  - cold 65
  - control 65
  - delete 65
  - lock 65
  - move 65
  - structure 65
  - type
    - client 65
    - private 65
    - private, managed by CQS 65
    - values 65
- Queue Control Facility (QCF)
  - cold queues
    - requeuing 222
  - structures
    - monitoring 222
    - requeuing 222
- queue manager
  - page fixing 412
  - trace report 430
- Queue Manager
  - in XRF environment 209
  - planning for 209
- queue types
  - in a shared-queues environment 206

## R

- RACF 175
  - AO application security 315
  - AO command security 319
  - command security 157
  - MSC
    - when IMS calls RACF 324
  - online change 337
  - resource classes
    - resource access security 328
- RACF (Resource Access Control Facility)
  - APPL class
    - example 158
  - authorization 175
  - implementing with IMS 332
  - IMSplex
    - authorization 158

- RACF (Resource Access Control Facility)
  - (continued)

- IMSplex (continued)
  - security 175
  - security example 175
- OPERCMD5
  - example 158
  - OM 158
- PassTickets 331
  - creating PassTickets 312
  - VTAM VGR environment 311
- password protection 314
- placement of data sets 644
- planning considerations 644
- protecting databases 339
- protecting system libraries and data sets 339
- resource class descriptor table 332
- resource classes for IMS 332, 342
- signon password reverifying 314
- RACF security
  - IMSplex checking 159
- RDS (restart data set)
  - allocating 110
  - as surveillance mechanism 611
  - content 98
  - definition 92, 597
  - establishing as surveillance 606
  - parameter
    - for surveillance 649
    - in member DFSHSBxx 606
  - placement of in XRF 656
- read access 237
- read-only access 237
- readiness level
  - database 672
- real storage
  - defining DFSFIXxx 611
  - need for during takeover 600
  - page fixing 611
  - used during XRF tracking 611
- reason codes
  - INITIATE OLC command 453
  - online change failure 459
- rebuilding structures 185
- recommendations
  - shared queues
    - using a resource structure 44
- RECON
  - as part of an IMSplex 21
  - automatic loss notification 51
  - data-sharing environment, in a 234
  - parallel access 143
- RECON data set
  - accessing in parallel 536
  - active 544
  - ALLOC record 554
  - allocation 535
  - and RSR tracking system 533
  - availability 536
  - avoiding deadlock situations 538
  - backing up before reorganization 559
  - BACKOUT record 548
  - backup 556
  - both sets are unusable 560
  - CA record 548
  - CAGRP record 548

- RECON data set (*continued*)
  - changing log control records 486
  - commands
    - BACKUP.RECON 556
  - concurrent image copy data set 520
  - contention problems 536
  - creating, recommendations when 542
  - data sharing record types 554
  - data-sharing records 499
  - Database Image Copy 2 data set 518
  - DBDSGRP record 549
  - defining 535
  - defining requirements in 489
  - description 535
  - DSECTs that map to records 555
  - dynamic allocation 541
  - enqueue problems, causes of 576
  - extending 542
  - GSG record 552
  - header records 546
  - HSSP image copy data set 523
  - IMAGE record 553
  - initial access 544
  - initialization 536
  - initializing 505
  - input/output error processing 560
  - log information 485
  - LOGALL record 554
  - loss notification 562
  - maintaining 556
  - maintaining records 556
  - monitoring in data sharing environment 254
  - OLDS information 94
  - output
    - time stamps 567
  - overriding DBRC security 565
  - overview 480, 535
  - parallel access 536
  - parallel access, planning for 537
  - planning considerations 535
  - preserving integrity of 563
  - processing
    - LSR (Local Shared Resources)
      - option 542
  - record mapping 555
  - record types 545
  - records 545
    - database 549
    - database data set 549
    - global service group 552
    - image copy 553
  - records, maintaining 556
  - RECOV record 554
  - recovering 560
  - recovery record types 548
    - ALLOC 554
    - BACKOUT 548
    - CA 548
    - CAGRP 548
    - DBDSGRP 549
    - GSG 552
    - IMAGE 553
    - LOGALL 554
    - RECOV 554
    - REORG 553
    - SSYS 554
- RECON data set (*continued*)
  - registering databases 507
  - REORG record 553
  - reorganization, backing up
    - before 559
  - reorganizing 559
  - reorganizing online 559
  - reorganizing procedure 559
  - replacing a discarded RECON 562
  - replacing damaged 560
  - retrieving log-related information
    - from 488
  - security considerations 544
  - serialization 538
  - SLDS information 96
  - spare 544
  - SSYS record 554
  - subsystem (SSYS) record 554
  - time stamps
    - output 567
  - tracking changes to 566
  - unusable 560
  - use of 98
  - VSAM CREATE mode 541
- RECON I/O exit routine
  - tracking changes to the RECON 566
- RECON initialization token (RIT) 497
- records
  - BACKOUT 548
  - change accumulation group 548
  - change accumulation run 548
  - data sharing 554
  - database 549
  - database data set 549
  - database data set group 549
  - database recovery 548
  - DBDS group 549
  - HALDB online reorganization 531
  - in RECON data set 545
  - log allocation 554
  - log data set 546
  - online reorganization 531
  - RECON
    - deleting unnecessary 556
    - maintaining 556
  - RECON data set header 546
  - reorganization 553
  - subsystem 554
- records restart 62
- RECORDSIZE keyword
  - DEFINE CLUSTER keywords 542
- RECOVER commands 494
- recoverable databases
  - CHANGE.DB command 516
  - Image Copy utilities 516
  - INIT.DB command 516
- recoverable service element (RSE)
  - definition 586
- recovering a database
  - archiving log records 487
  - batch support 509
- recovering data in message queues and databases
  - XRF 615
- recovering the RECON data set
  - RECON loss notification 562
- recovery
  - backward
    - definition 115
  - batch backout 267
  - complications 117
  - CQS functions 67
  - data sharing 268
    - IRLM 269
    - planning procedures 266
    - sysplex 283
    - without DBRC 511
  - database
    - DBRC role 493
    - dynamic backout with data sharing 509
    - forward recovery with data sharing 509
    - setting up recovery mechanisms 509
    - without DBRC 511
  - DBCTL environment 113
  - DBRC
    - introduction 113
  - description of 89
  - determining extent for RSR 672
  - dynamic backout 267
  - example of system without recovery mechanisms 89
  - executing related functions 350
  - facilities 266
  - forward
    - data sharing 268
    - definition 116
  - groups 549
  - in an IMSplex 23
  - in data sharing
    - planning procedures 509
    - transaction recovery 269
  - information 62
  - introduction to 89
  - maintaining records for OLR 533
  - mechanisms 91
  - mechanisms, setting up 509
  - overhead of 118
  - period, of image copy data sets 521
  - planning 266
  - point-in-time 490
  - procedures, reviewing 82
  - process 90
  - RECON data set 560
  - record (RECOV) 554
  - recovering CQS 185
  - RSR (Remote Site Recovery)
    - overview 663
  - RSR processing 672
  - services
    - introduction 113
  - steps 90
  - structures 187
  - sysplex
    - data sharing 283
  - system requirements 90
  - time stamp 490
  - using logs 92
  - utilities
    - input to 97
    - introduction 113

- recovery (*continued*)
  - utilities, input to 96
  - what IMS cannot recover 118
  - XRF environment 114
- recovery control data set 535
- I/O error processing 560
- recovery level tracking (RLT) 672
- recovery log data set (RLDS)
  - condensing 102
  - copying 101
  - creation 101
  - defining 110
  - definition 92
  - log 97
  - records 546
- recovery period
  - definition 521
- recovery records, database 548
- recovery utilities
  - database 494
- recovery, database
  - process overview 489
  - recovery, database
    - concepts 489
- RECOVPD parameter
  - usage 521
  - using 521
- Region and Jobname Report
  - IMS Monitor
    - DB/DC 737
    - DCCTL 789
- Region IWAIT report 416
- region occupancy
  - measuring 743, 795
- Region Summary report
  - for region occupancy percentage 414
  - NOT-IWAIT time implications 375
- Region Summary Report
  - IMS Monitor
    - DB/DC 740
    - DCCTL 791
  - IMS Monitor (DBCTL) 772
- Region Wait Report
  - IMS Monitor
    - DB/DC 740
    - DBCTL 772
    - DCCTL 791
- registering area data sets
  - in the RECON data set 507
- registering databases
  - in the RECON data set 507
- registering interest
  - in a shared queue
    - concept 206
  - in transactions 207
  - shared queues
    - in LTERMs 207
    - in MSC (multiple systems
      - coupling) resources 208
- registration, authorizing 178
- relative importance
  - assigning to address space 372
  - Workload Manager 372
- remote processing
  - MSC (multiple systems
    - coupling) 205
  - shared queues 205
- Remote Site Recovery (RSR)
  - components of 665
  - data sharing support 682
  - DBRC support of 501
  - DL/I database tracker 668
  - error handling at active site 698
  - error handling at remote site 700
  - example 687
  - extent of recovery, determining 672
  - Fast Path database tracker 669
  - hardware replication 691
  - ILS (isolated log sender) 695
  - IMS and DB2 for z/OS 688
  - IMS workload on multiple z/OS
    - images 695
  - IMSplex tracking 685
  - in an IMSplex 23
  - initializing
    - active site 696
    - tracking site 697
  - installing 691
  - introduction 663
  - isolated log sender (ILS) 667, 695
  - log management 685
  - log router 666
  - multiple z/OS images, running IMS
    - on 695
  - naming conventions for 669
  - online reorganization 533
  - overview 663
  - processing 672
  - requirements 665
  - security
    - establishing 703
    - terminals 703
  - shared queues, planning for 231
  - software replication 692
  - system definition requirements 674, 696
  - takeover definition 672
  - terminal security 703
  - tracking an IMSplex 685
  - Transport Manager Subsystem (TMS) 666
  - XRF support 681
  - XRF, comparison with 680
- Remote Site Recovery (RSR) overview 114
- remote takeover
  - restarting active IMS systems 684
- remote takeover, definition 672
- Removing an SSYS from DB
  - Authorization 572
- removing contention
  - I/O resource
    - tuning 428
- REORG record
  - fields used in online
    - reorganization 531
  - RECON data set 553
- reorganization record 553
- reorganizing databases 259
- reorganizing RECON
  - procedure 559
  - using CHANGE.RECON 559
- report
  - job summary 380
- report (*continued*)
  - system 380
  - tracing 380
- report output sequence
  - IMS Monitor Report
    - DBCTL 769
- reports
  - Call Summary 362
  - IMS Monitor 378
  - IMSASAP II 362
  - Statistical Analysis utility 349
- reports, IMS
  - using 709
- REPRO command
  - restoring RECON data sets 560
  - using for backup 556
- requirements
  - analyzing 81
  - CQS
    - hardware 62
    - software 62
  - IMS system 81
- RESERVE command
  - during backup 556
- RESERVE macro 538
- resetting the GENMAX parameter 569
- resource
  - changing online 437
  - monitoring in data sharing
    - environment 252
  - structure
    - changes logged 177
    - recovery 187
  - structures 66
- resource access security (RAS)
  - dependent regions 328
  - RACF resource classes 328
- resource classes
  - correspondence to IMS definition 332
  - RACF
    - resource access security 328
- resource cleanup module (DFS MRCL0)
  - installing 132
  - uninstalling 133
- Resource Management Facility II
  - as monitoring tool 375
- Resource Management Facility II (RMF II)
  - paging rates, using for 375
- Resource Manager
  - online change
    - restrictions when RM not used 49
- Resource Manager (RM)
  - administration tasks 169
  - configuration requirements 54
  - configuring CSL without 54
  - function provided 44
  - global status
    - processing errors 170
  - maintaining global command
    - status 170
  - maintaining global resource
    - information 169
  - overview 44
  - PLEXPARM= parameter 170
  - resource information stored
    - for databases 171
    - for DEDB areas 172



- Resource Manager (RM) *(continued)*
  - resource information stored *(continued)*
    - for transactions 172
  - resource structure 44
    - examples of global resources 169
    - repopulating 174
    - storing resource information 169
  - resource name uniqueness
    - disabling enforcement 299
  - resource structure 51, 169
    - configuration 53
    - CQS 174
    - defining 215
    - failure 174
    - IMSplex
      - message destination resource 172
      - Transaction Manager
        - resources 172
    - information stored 169
    - online change
      - restrictions when resource structure 49
    - rebuild 174
    - recovery 173
    - resource information stored
      - for databases 171
      - for DEDB areas 172
      - for transactions 172
    - serialized programs 169
  - resource type consistency
    - disabling enforcement 300
  - resource utilization
    - changing 402
  - resources
    - IMSplex, in an
      - callable services 300
      - managing 299
    - modifying 351
    - monitoring in data sharing
      - environment 352
    - name uniqueness 299
      - LTERM 301
      - user names 303
    - VTAM terminal node 302
    - TM resources
      - MSNAME 301
      - transactions 302
      - user IDs 304
      - user names 303
      - VTAM terminal nodes 302
    - type consistency 300
- response time
  - as a user oriented objective 370
  - criteria definition 370
  - XRF takeover 632
- restart
  - automatic
    - definition 117
  - cleaning up after failed attempt 464
  - data sharing 268, 509
  - emergency
    - definition 117
  - normal
    - definition 117
  - overview 117
  - use of logs 74
- restart data set 597
- restart data set (RDS)
  - allocating 110
  - content 98
  - definition 92
- restart in data sharing
  - after an IMS failure 512
  - after DBRC failure 512
  - emergency, after DBRC failure 512
  - emergency, after IMS failure 512
  - normal 512
- restarting IMS 355
- restrictions
  - concurrent image copy (CIC) 519
  - online reorganization 533
- return codes
  - binder
    - interpreting 203
  - INITIATE OLC command 453
  - online change failures 459
- returning failed active IMS as new active IMS 621
- REUSE parameter
  - image copy data sets for future use 520
- REUSE parameter commands
  - INIT.DBDS 522
- reusing image copy data sets 522
- reviews
  - design 82
- REXX SPOC API 48
- RIT (RECON initialization token) 497
- RLDS (recovery log data set)
  - accumulating changes using DFSUCUM0 525
  - condensing 502
  - copying 101
  - creation 101
  - defining 110
  - definition 92
  - log 97
  - records 546
- RM 39
- RM (Resource Manager)
  - configuration requirements 54
  - configuring CSL without 54
  - defining 25
  - function provided 44
  - overview 44
  - resource structure 44
  - terminals, managing 51
  - XRF 585
- RMF II (Resource Management Facility II)
  - as monitoring tool 375
  - for I/O analysis 434
  - paging rates, using for 375
  - used in tuning 407
- RMGENJCL (/RMGENJCL) command
  - generating JCL 493
- RMxxxxxx (/RMxxxxxx) commands
  - overview 482
- ROLB call
  - dynamic backout 509
- routing codes
  - 403
  - with online change 403
- RRS
  - Archive Log Stream
    - considerations 137
  - ODBM support 148
- RRSAF (DB2 Recoverable Resource Services attachment facility) 13
- RSE (recoverable service element)
  - definition 586, 597
  - notifying the availability manager (AVM) 649
  - RSENAME keyword 649
- RSENAME= keyword for XRF 649
- RSR
  - DRD 684
  - initializing 696
  - overview 663
  - system failure
    - resume service 585
- RSR (Remote Site Recovery)
  - active IMS 665
  - alternate IMS system 665
  - components of 665
  - data sharing support 682
  - DB/DC environment 11
  - DBRC support of 501
  - DL/I database tracker 668
  - error handling at active site 698
  - error handling at remote site 700
  - example 687
  - extent of recovery, determining 672
  - Fast Path database tracker 669
  - global online change, and 142
  - hardware replication 691
  - IMS and DB2 for z/OS 688
  - IMS workload on multiple z/OS images 695
  - IMSplex tracking 685
  - initializing
    - active site 696
    - tracking site 697
  - installing 691
  - introduction 663
  - isolated log sender (ILS) 667
  - log management 685
  - log router 666
  - multiple z/OS images, running IMS on 695
  - naming conventions for 669
  - online reorganization 533
  - overview 663
  - processing 672
  - requirements 665
  - security
    - establishing 703
    - terminals 703
  - shared queues, planning for 231
  - software replication 692
  - system definition requirements 674, 696
  - takeover definition 672
  - terminal security 703
  - tracking an IMSplex 685
  - Transport Manager Subsystem (TMS) 666
  - XRF support 681
  - XRF, comparison with 680

- RTCODE macro statement
  - used with online change 403
- Run Profile Report
  - adding generalized processing ratios
    - DB/DC 737
    - DBCTL 770
    - DCCTL 789
  - IMS Monitor
    - DB/DC 737
    - DBCTL 770
    - DCCTL 789
  - overview 737
- runtime resource definitions
  - changing 447

## S

- sample exit routine
  - SMP/E assemble 204
- SB (OSAM Sequential Buffering) 416, 419
- SB-Summary Report 814
- scenarios
  - online change 457
- scheduling
  - AOI transactions 230
  - application programs
    - against unavailable data 71
- scheduling algorithm
  - factors in region occupancy 414
- scheduling for critical transactions
  - examining 829
- SCI 39
- SCI (Structured Call Interface)
  - automatic RECON loss
    - notification 51
  - configuration requirements 54
  - functions provided 45
  - initializing 25
  - overview 45
  - registration exit routine 51
  - security 175
  - XRF 585
- secondary terminals
  - XRF 647
- security 307
  - activating IMS security 331
    - defining the SECURITY macro 331, 344
    - initializing RACF 332, 344
  - AO (automated operator) application programs 74, 315
  - changing online 442, 446
  - choices made during system
    - definition 308
  - CMD call 315
  - commands allowed for RECON 563
  - CPI-C driven application program
    - considerations 321
  - database
    - RACF security 339
    - segment and field-level sensitivity 339
  - DB/DC
    - overview 307
  - DBCTL 342
  - DBCTL considerations 341
- security (*continued*)
  - DCCTL
    - overview 307
  - defining
    - EXEC statement parameters 335
  - dependent regions
    - no signon 329
    - resource access security (RAS) 328
  - design considerations
    - authorizing commands 312
    - authorizing transactions 312
    - choosing types of security 310
    - DBCTL environment 343
    - limiting access from a dependent region 327
    - limiting access from a terminal 310
    - master terminal 315
    - using RACF 314
  - display bypass and password
    - masking 338
  - encryption
    - using cryptographic support 341
    - using the ICSF/CCA interface 341
    - using the Segment Edit/Compression routine 341
    - VTAM terminals 341
  - ETO 327
  - facility class 175
  - Fast Path considerations
    - DB/DC 321
    - in DBCTL 344
  - ICMD call 315
  - implementing online 442, 446
  - IMS Connect 330
  - IMSplex 30
    - RACF OPERCMDS class 158
  - JMP applications 321
  - MSC
    - MSCSEC= 324
    - transactions 324
    - when IMS calls DFSCTRNO 324
    - when RACF and exit routines are called 324
  - ODBA application program
    - considerations 344
  - ODBA application programs 322
  - ODBM (Open Database Manager) 150
  - online changes in DBCTL 347
  - Open Database Manager (ODBM) 150
  - Open Transaction Manager Access (OTMA) 330
  - OTMA (Open Transaction Manager Access) 330
  - password verification 311
  - physical 338
  - preparing exit routines 331
  - RACF 175
  - RACF (Resource Access Control Facility) 442, 446
  - Remote Site Recovery (RSR)
    - terminal security 703
- security (*continued*)
  - Resource Access Control Facility (RACF) 442, 446
  - resources that can be protected 307, 341
  - scope in MSC and shared queues environments 322
  - signon verification 311
  - startup options 335
  - Structured Call Interface 175
  - system libraries and data sets
    - DLISAS procedure 339
    - IMS procedure 339
    - overview 339
  - system startup options 346
  - TCO (Time-Controlled Operations) 320
  - terminal, default 308
  - user
    - no signon 329
  - security considerations
    - RECON data set 544
  - security facilities
    - DB/DC resources 309
    - DCCTL resources 309
  - SECURITY macro
    - defining 331
    - defining for DBCTL 344
    - RACFCOM keyword 331
    - RACFTERM keyword 331
    - RCLASS keyword 331
    - SECCNT keyword 331
    - SECLVL keyword 331
    - TERMNL keyword 331
    - TYPE keyword 331
  - SECURITY macro statement
    - DBCTL environment 342
    - overriding values with JCL 346
    - SECLVL keyword 315
    - TERMNL keyword 315
  - security options
    - for dependent regions 70
  - security violations 337
    - notifying the master terminal 337
    - recorded on system log 337
    - setting a threshold 337
  - Segment Edit/Compression exit routine 341
  - segment-level sensitivity 339
  - Sequential Buffering Detail report 816
  - Sequential Buffering, OSAM 416, 419
  - serial applications
    - processing
      - in a shared-queues environment 210
  - serial programs
    - GBL\_SERIAL\_PGM 173
    - managing serialization in a sysplex environment 173
    - sysplex, managing in a 173
  - serial transactions
    - processing
      - in a shared-queues environment 210
  - serialization
    - of RECON data set 538

- service
  - considerations 191
  - corrective service
    - Authorized Program Analysis Report (APAR) 191
  - installing 195
  - description 191
  - installing 195, 200
    - attention notice 195
  - preventive service
    - installing 195
    - program temporary fix (PTF) 191
  - process 191
  - Recommended Service Upgrade (RSU) 192
  - special considerations
    - IVP 202
    - non-SYSDEF target libraries 202
  - SYSMOD Packaging 191
- service class 373
- service definition 373
- service group (SG) 669
- session recovery
  - definition, XRF 597
  - speed of after XRF takeover 616
- session-switching at takeover
  - backup session 616
- session-switching at XRF takeover
  - class-1 terminals 597
- SETO (set options) call 469
  - allocation error 472
  - controlling print data sets 471
  - descriptors, Spool API 471
  - express alternate PCB 472
  - ISRT call, writing data and 470
  - output DD statement 470
  - purge call (PURG) 470
  - ROLB call 471
  - ROLL call 471
  - ROLS call 471
  - SETS call 471
  - SETU call 471
  - writing data 470
  - XRF, and 472
- SG (service group) 669
- shadow database, definition 665
- shared data 233
- shared queues
  - autologon terminal 209
  - cloned configuration 211
  - concepts 205
    - back-end IMS 205
    - Common Queue Server (CQS) 205
    - front-end IMS 205
    - queue types 206
    - when IMS registers and deregisters interest in transactions 207
  - configuring 211
  - conversational transactions
    - concepts 211
  - CQS
    - execution parameters 218
    - global structure definition 218
    - initialization parameters 218
    - local structure definition 218
  - defining CFRM policy 215
- shared queues (*continued*)
  - example 215
  - defining IMS control region execution parameters 218
  - enabling 215
  - environment
    - AOI transactions 230
    - APPC messages 228
    - OTMA messages 228
    - performance planning 426
    - planning for MSC (multiple systems coupling) 230
  - fast path message queue list structures
    - example 215
  - Fast Path message queue list structures
    - defining 215
  - locking messages on queue
    - concept 208
  - message flow
    - concepts 208
  - message queue list structures
    - defining 215
    - example 215
  - message switching
    - Destination Creation exit routine (DFSINSX0) 209
  - name uniqueness 299
  - partitioned configuration 212
  - planning for 205
  - program-to-program switching
    - Destination Creation exit routine (DFSINSX0) 209
  - registering and deregistering interest in LTERMs
    - concept 207
  - registering interest
    - concept 206
  - registering interest in MSC (multiple systems coupling) resources
    - concept 208
  - remote processing
    - defined 205
  - resource structure
    - defining 215
  - serial application processing 210
  - serial transactions processing 210
  - TM resources
    - managing 299
  - tuning performance 232
    - parameters 232
    - structures 232
  - undefined transactions
    - Destination Creation exit routine (DFSINSX0) 209
  - without a CSL 54
  - z/OS log stream
    - defining 217
- shared secondary index database
  - status as global resource 170
- SHARELVL parameter
  - specification descriptions 499
- SHAREOPTIONS keyword
  - DEFINE CLUSTER keywords 542
- sharing level
  - assigning one with DBRC 499
  - DBRC, assigning in 499
- sharing level (*continued*)
  - definitions of values 499
- shutdown
  - commands 116
  - IMS 357
  - IMSplex, in an 304
  - overview 116
- SIGN (/SIGN) ON/OFF Security exit routine 331
- Sign-on exit routine (DFSSGNX0)
  - using with ETO 327
- signon verification
  - DFSDCxxx PROCLIB member 311
  - exit routine 331
  - for all static terminals 311
  - RACF PassTickets 331
  - security 311
  - VTAM terminals 311
  - with RACF 314
- simulation of online execution 364
- single logging 106
- single point of control (SPOC)
  - functions 45
  - IMS Control Center 45
  - REXX 48
  - TSO SPOC 45
  - user-written 45
- size calculation for SSYS record 572
- SLDS (system log data set)
  - accumulating changes using DFSUCUM0 525
  - allocating 96
  - block size 110
  - condensing 102
  - copying 101
  - creating 96
  - defining 110
  - definition 92
  - environment 96
  - log 96
  - RECON data set 96
  - records 546
  - single or dual 110
- SLDS stop time, locating the last in the RECON data set 567
- SMF
  - DDCONS parameter considerations 137
- SMP/E
  - commands
    - ACCEPT 196, 200
    - ACCEPT CHECK GROUPEXTEND BYPASS(APPLYCHECK) 200
    - ACCEPT GROUPEXTEND 195
    - ACCEPT GROUPEXTEND BYPASS(APPLYCHECK) 196, 200
    - APPLY 196
    - APPLY CHECK GROUPEXTEND 195, 200
    - APPLY GROUPEXTEND 195, 200
    - attention notice 200
    - CLEANUP 196
    - GENERATE 196, 202
    - JCLIN 196
    - LIST 196
    - RECEIVE 195, 196, 200



- SMP/E (*continued*)
  - commands (*continued*)
    - RESTORE 196, 200
    - UNLOAD 196
    - ZONEDELETE 196
    - ZONEMERGE 196
  - installation methods 195
  - service 191
- SNA terminals
  - XRF 595
- SNA terminals in XRF 588
- SNAPQ checkpoint 610
- software requirements
  - CQS 62
- space requirements, RECON data
  - set 535
- SPEED keyword
  - DEFINE CLUSTER keywords 542
- SPOC (single point of control)
  - functions 45
  - MTO 45
  - OM API 152
  - overview 45
  - requirements 45
  - REXX SPOC API 45
  - TSO SPOC 45
  - user-written 45
  - WTOR 45
- Spool API 465
- SSM EXEC parameter
  - connecting to external subsystems 13
- staging libraries 438
- start
  - cold
    - performing 463
  - data sharing processing 249
  - IRLM subsystem 249
- START (/START) AVM command 608
- START (/START) command
  - defining sharing level with 499
- START (/START) IMS command 608
- START (/START) SUBSYSTEM SSM command 13
- START (/START) SURVEILLANCE
  - command
    - XRF usage 597
- starting IMS 355
- STARTNEW parameter
  - usage 560
- startup parameters 218
- Statistical Analysis utility 349
- Statistical Analysis utility (DFSIST50)
  - Application Accounting report 832
  - calculating transaction loads 826
  - execution savings 825
  - for message traffic 416
  - input 825
  - Line and Terminal report 826
  - Messages report 828
  - multiple systems, using with 369
  - reports produced, descriptions, and examples 825
  - transaction profiles 832
  - Transaction Response report 826
  - use of the IMS system log 375
  - used in tuning 407
  - utilization for accounting 832
- statistics
  - transaction-level
    - analyzing 388
    - interpreting 388
    - monitoring 388
- stop
  - IRLM subsystem 250
  - OLDS 94
- STOP (/STOP) SURVEILLANCE
  - command
    - XRF usage 597
- structure
  - alter 179
  - authorizing connections to 178
  - checkpoint, initiating 181
  - copy 188
  - duplexing 66, 188
  - enabling 188
  - full, monitoring 183
  - functions 66
  - overflow 184
    - function 183
  - structure full monitoring 184
  - overflow processing 66
  - pair 65
  - rebuild 66
    - initiating 186
  - recovery 187
  - repopulation 186
  - size 179
  - types 65
- structure full
  - managing 182
- structure usage
  - managing 182
- Structured Call Interface
  - security 175
- Structured Call Interface (SCI)
  - administration 175
  - configuration requirements 54
  - functions provided 45
  - overview 45
- structures
  - resource 66
- subsystem (SSYS) record
  - deleting a 572
  - initializing during IMS restart 554
  - size calculation 572
  - working with 572
- subsystems
  - connecting 353
  - disconnecting 353
  - supervisor call 130
- SURV parameter in member
  - DFSHSBxx 606
- surveillance
  - of the active IMS 611
  - XRF options 649
- surveillance mechanisms
  - absence of signal causing takeover 613
  - absence of signal causing XRF takeover 597
  - changing, XRF 597
  - definition for XRF 597
  - establishing for XRF 597
  - starting, XRF 597
- surveillance mechanisms (*continued*)
  - stopping, XRF 597
  - XRF example 597
- surveillance mechanisms for XRF
  - establishing 606
  - interval value, setting 606
  - summary of parameters 606
  - timeout value, setting 606
- SVC
  - type-2, binding 131
  - type-2, loading 131
- SVC (supervisor call) modules 130
- SWITCH (/SWITCH) SYSTEM command
  - planned takeover 613, 620
  - takeover process 613, 618
- SWITCH parameter in DFSHSBxx 606
- switching devices
  - XRF 656
- switching-priority of sessions
  - class-1 terminals, defining 648
  - default value 646
- sync point 91
  - application program 92
  - definition 91
  - overview 91
- synchronization phase
  - description of 610
  - illustration of 610
  - using SNAPQ checkpoint 610
- synchronization point
  - definition 76
  - in transaction flow 384
- Syntax Checker
  - starting with IMS Application Menu 41
- SYS1.NUCLEUS
  - attention notice 131
  - discussion of 131
  - installing IMS type-2 SVC 131
- SYSIN/SYSOUT in program testing 366
- SYSLIB
  - proper concatenation 203
- sysplex
  - data sharing 283
  - data sharing environment 270
  - local online change 450
  - online change 449
  - programs, serialized 173
  - serialized program management 173
- sysplex data sharing 271
  - buffer invalidation 271
  - calculating size of structures 280
  - concepts 271
  - configurations 274
  - converting batch jobs to BMP jobs 279
  - coupling facility 273
  - structures
    - changing size 283
  - terminology 271
  - when to use 279
  - XRF 274
- sysplex environment
  - APPC messages 228
  - asynchronous 228
  - synchronous 228
  - eligibility for processing 227

- sysplex environment (*continued*)
  - messages 227
  - OTMA messages 228
    - asynchronous 228
    - synchronous 228
- system
  - administration
    - overview 3
  - changing system design 395
  - data set
    - backup 112
  - documenting the system 85
  - modifying system design 395
  - monitoring 369
  - performance
    - overview 408
  - report 380
  - testing
    - establishing a test system 360
    - setting up a test system 360
- system administration
  - considerations 125
  - tasks 125
- system checkpoint 91
  - definition 75
- system checkpoint data set, CQS 98
- system checkpoint, initiating 180
- System Configuration Report
  - IMS Monitor
    - DB/DC 737
    - DBCTL 770
    - DCCTL 789
- system console organization
  - XRF 595
- system data sets, protection 339
- system definition
  - changing online 441, 445
  - documentation 85
  - modifying for online change 402
  - requirements for data sharing 246
  - simplified IMSplex 25
  - type ALL
    - when to perform 196
  - XRF complex 645
- system definition changes
  - controlling 399
  - coordinating security definitions 401
  - for online change 402
  - selecting a definition type 400
- system definition macro statements
  - coding for XRF 646
  - reprocessing requirements 400
- system design changes, effect of 395
  - databases modified 395
  - exit routines 395
  - for online change 395
  - message format changed 395
  - network control 395
  - output changed 395
  - programs modified 395
  - scheduling changes 395
  - security maintenance 395
  - terminal attachment 395
  - transactions modified 395
  - tuning changes 395

- system failure
  - resuming service
    - IMSplex 585
    - RSR 585
    - XRF 585
  - synchronizing IMS restart 287
- system library, protection 339
- system log
  - introduction 74
  - security violation records 337
- system log data set (SLDS)
  - allocating 96
  - block size 110
  - condensing 102
  - copying 101
  - creating 96
  - defining 110
  - definition 92
  - environment 96
  - log 96
  - RECON data set 96
  - records 546
  - single or dual 110
- system messages
  - at initialization 610
  - at XRF takeover 591, 638
  - during I/O prevention 600
  - process triggered by 615
  - when IRLM operation resumes 619
  - when user logs on to alternate IMS 601
- system programmers
  - understanding I/O prevention 600
- system recovery
  - after failure 287
- system resource manager (SRM) 632
  - contribution to XRF 600
- system startup
  - JCL security options 335, 346
  - security 335
- System Support Program (SSP)
  - generating control blocks for backup sessions 643
  - requirement for XRF 595
  - XRF process, contribution to 605
- system-managed rebuild 185
- system, as part of RSR name 669
- systems management tasks 39

## T

- tailoring
  - execution JCL for data sharing 248
  - IMS for XRF 649
  - z/OS 248
- takeover
  - as users see it 638
  - conditions for XRF 591
  - definition 672
  - definition for XRF 586
  - differences in takeover process 618
  - effect on non-XRF jobs on alternate IMS system 632
  - effect on terminals 616
  - establishing conditions for XRF 591
  - IMS processing for XRF 597
  - messages 638

- takeover (*continued*)
  - phase, different processes 618
  - planned
    - process 614, 620
    - uses of 591
  - problem determination after 591
  - setting criteria for 606
  - system messages during, XRF 597
  - VTAM forced takeover for XRF 601
  - XRF, as users see it 588
  - XRF, example of 591
  - XRF, uses of 620
- takeover conditions
  - ISC link fails to send signals 606
  - log records fail to appear 606
  - planned 613
  - RDS signals fail to appear 606
  - XRF 613
- takeover phase
  - XRF, description of 614
- tasks
  - system administration 125
  - tasks performed with DBRC 481
- TCO (Time-Controlled Operations)
  - securing
    - commands from 320
    - loading of 320
  - security 320
- Teleprocessing Network Simulator (TPNS) 367
- terminal
  - autologon
    - shared queues 209
  - session balancing
    - persisting affinity 294
- terminal failure
  - XRF limitations 591
- TERMINAL macro statement
  - BACKUP keyword 648
  - NAME keyword 647, 648
- terminal profiles, documentation 85
- terminal status
  - resetting 294
- terminals
  - as a TM resource in a sysplex 302
  - defining priority of session
    - recovery 648
  - IMSplex, in an 302
    - managing 299
  - logon in the XRF complex 601
  - managing globally 51
  - name uniqueness 299
  - Remote Site Recovery (RSR)
    - security 703
    - requirement for XRF 595
  - RM definition 302
  - system definition macro
    - keywords 648
  - VTAM terminal nodes 302
- TERMINATE OLC command 445, 456
- termination phase
  - description of 623
- test cycle
  - maintenance 194
- test environments
  - monitoring 362

- testing
  - aids 364
  - database 361
  - ensuring network readiness 362
  - establishing a test system 360
  - MFS formats online 364
  - MFSTEST mode 364
  - monitoring in test environments 362
  - online 363
  - online system 359
  - operational procedures 361
  - performance 367
  - phases in 359
  - procedures, operational 361
  - programs 366
  - QCF, with 367
  - setting up test system 360
  - simulating online execution 364
  - stress 367
  - SYSIN/SYSOUT 366
  - system
    - setting up 360
  - test database 361
  - test system 359
  - with Batch Terminal Simulator 364
  - with online change 365
- time sharing option (TSO)
  - running on alternate IMS system 632
- time stamp
  - recovery 490
- Time-Controlled Operations (TCO)
  - security 320
- timeout value
  - setting 606
- timeout value, definition of 607
- Tivoli NetView for z/OS 606
  - VTAM application name
    - changing in USERVAR table 616
    - communicating new name 606
  - XRF process, contribution to 606
- TM batch
  - DBBBATCH 18
  - DLIBATCH 18
  - GPSBs, specifying 18
  - overview 18
- TM resources 300
  - APPC Descriptors 300
  - LTERM (logical terminal) 301
  - MSNAME 301
  - transactions 302
  - user IDs 304
  - user names 303
  - VTAM terminal nodes 302
- TMS (Transport Manager Subsystem) 666
- tools
  - productivity 113
- TPEND exit
  - VTAM failures 597
- TPNS (Teleprocessing Network Simulator) 367
- trace
  - CTRACE records 242
  - GTF (Generalized Trace Facility) 375, 392
  - GTF Detail Trace report 375
- trace (*continued*)
  - GTFPARS Job Summary and Detail
    - Trace report 430, 433, 434
  - IMS Monitor 362, 375, 392
  - IMSPA DC Queue Transaction report 430
  - IRLM activity 242
  - when to avoid using 423
  - TRACE (/TRACE) command
    - for program isolation 381
    - IMS Monitor 378
    - to turn on IMS Monitor 388
  - TRACE (/TRACE) SET OFF
    - command 32
  - TRACE (/TRACE) SET ON
    - command 32
  - TRACE CT command 380
  - trace facility 382
  - trace report 380
  - tracing 382
    - BPE components 35
    - CTRACE, using 380
    - defining external data sets 36
    - GTF trace 380
    - log 101
    - program isolation and lock 381
    - starting BPE external tracing 38
    - stopping BPE external tracing 38
    - writing to BPE internal trace tables 35
    - writing to external data sets 36
  - tracking IMS
    - definition 665
  - tracking phase
    - description of 611
    - illustration of 613
  - trademarks 847
  - TRANSACT macro
    - AO command security 317
  - TRANSACT macro statement
    - used with online change 403
  - transaction
    - authorization 312
    - definition 6
    - exit routine preparation 331
    - Fast Path
      - potential 73
    - flow and IMS Monitor events 733
    - flow of events 384
    - for program-to-program switch 312
    - RACF PassTickets 331
    - security with no user signon 329
    - with exit routine or RACF 312
  - Transaction Manager
    - resources 300
      - APPC Descriptors 300
      - LTERM (logical terminal) 301
      - MSNAME 301
      - transactions 302
      - user IDs 304
      - user names 303
      - VTAM terminal nodes 302
  - transaction profiles
    - in capacity planning 405
    - obtaining base profiles 370
    - significant elements 370
- Transaction Queuing report
  - DBCTL
    - example 779
  - IMS Monitor (DCCTL) 802
- Transaction report 826
- Transaction Response report
  - description 826
  - example 826
- Transactional VSAM (TVS)
  - enabling 537
  - parallel access to RECON
    - enabling 537
- transactions
  - as a TM resource in a sysplex 302
  - auditing 828
  - changing online 442, 446
  - conversational
    - in a shared-queues environment 211
  - defining
    - type-1 AO commands 316
  - defining the user ID
    - type-2 AO commands 317
  - requeuing suspended 230
  - resource information stored 172
  - scheduling critical 829
  - security
    - MSC 324
  - serial
    - in a shared-queues environment 210
  - undefined to inputting IMS
    - Destination Creation exit routine (DFSINSX0) 209
  - when IMS registers and deregisters interest 207
- Transport Manager Subsystem (TMS) 666
- TSO SPOC 44
  - overview 45
  - starting with IMS Application Menu 41
- TSO, running on alternate IMS system 588
- tuning
  - ACBs into 64-bit storage 412
  - application control block placement 412
  - applications 425
  - as an iterative process 407
  - assessing I/O contention 434
  - defining a strategy 426
  - detecting processor resource problems 428
  - examining paging rates 412
  - implementation plan 401
  - overview 407
  - the IMS system 434
  - utilities 407
- TVS (Transactional VSAM)
  - enabling 537
  - parallel access to RECON
    - enabling 537
- TYPE macro statement, BACKUP
  - keyword 648
- type-1 AO application programs
  - security 315

- type-2 AO application programs
  - security 315
- type-2 command
  - RACF security
    - resource name and authorization 159
- type-2 command environment
  - defining 25
  - restrictions for global online change 49
- type-2 commands 38
  - configuration requirements 68
  - environment 68
  - requirements 68
  - type-2 command environment
    - restrictions for online change without RM 49
  - used without Resource Manager 41
- type-2 SVC
  - binding 131
  - loading 131

## U

- UCF (Utility Control Facility) with data sharing 511
- UIC, NOTIFY.UIC
  - updating the RECON data set 515
- unit of work (UOW)
  - tracking
    - shared queues environment 208
- UNLOCK (/UNLOCK) SYSTEM command
  - process during takeover 615
- UNLOCK SYSTEM command
  - ensuring database integrity 600
- UOW (unit of work)
  - tracking
    - shared-queues environment 208
- update access 237
- update capability
  - transferring 262
- updating control blocks in the alternate IMS system 611
- updating system definition
  - for online change 402
  - with multiple macro changes 399
- usage scenarios
  - online change 457
- user
  - IMSplex, in an
    - RM definition 303
- user IDs
  - as a TM resource in a sysplex 304
  - IMSplex, in an 304
    - name uniqueness 304
- user names
  - as a TM resource in a sysplex 303
  - IMSplex, in an 303
    - name uniqueness 303
- user-managed rebuild 186
- user, security when no signon 329
- usermods, service 191
- USERVAR
  - table
    - definition 601
    - example of 601

- USERVAR (*continued*)
  - table (*continued*)
    - IMS changing entry 616
    - procedures for initializing 601
    - variable definition 601
- USERVAR parameter
  - used for XRF 646
- USERVAR= keyword for XRF 649
- using DBRC, considerations for 513
- using logs
  - overview of 92
- utilities
  - Database Recovery Control utility (DFSURDB0) 509
  - Fast Path Log Analysis utility (DBFULTA0) 349
  - file select 177
  - for tuning 407
  - formatting print 177
  - interpreting
    - IMS Monitor Reports for DBCTL 767
    - IMS Monitor Reports for DCCTL 787
  - Interpreting
    - //DFSSTAT reports 811
  - interpreting DB-Monitor Reports 709
  - interpreting Statistical Analysis and Log Transaction reports 825
  - VSAM AMS 560
- utility
  - Database Change Accumulation 103
  - Global Online Change 463
- Utility Control Facility (UCF)
  - DBRC, relationship to 511
- utility control statement
  - INIT.CA 528
  - INIT.CAGRP 528

## V

- valid log subset, in data sharing to compress the size 526
- VGRS parameter
  - specifying a generic resource name 291
- violation control 347
- Virtual Storage Access Method (VSAM)
  - cache structures
    - changing size of 283
  - structure
    - calculating size 282
- virtual storage used during tracking 611
- Virtual Telecommunications Access Method 139
  - XRF requirements 595
- VS Pascal subroutines, preloading 422
- VSAM
  - data set definitions 249
- VSAM (Virtual Storage Access Method)
  - buffer pools 416, 419
  - cache structures
    - changing size of 283
  - structure
    - calculating size 282
- VSAM (Virtual Storage Access Method)
  - create mode
    - RECON data set 541
- VSAM AMS (access method services)
  - RECON data sets, restoring 560
  - restoring RECON data sets 560
- VSAM Buffer Pool Report
  - //DFSSTAT 812
  - description 812
  - fields 812
  - IMS Monitor (DB/DC) 753
  - IMS Monitor (DBCTL) 780
- VSAM-Buffer-Pool report
  - description 709
  - fields in the report 709
  - using the report 709
- VSAM-Statistics report
  - application program 715
  - description 715
  - fields in the report 715
  - I/O operations 715
  - improve performance
    - tuning the application program 715
    - tuning the database 715
- VTAM
  - as a TM resource in a sysplex 302
  - generic resource 295
  - generic resources
    - planning for 289
  - terminal nodes 302
- VTAM (Virtual Telecommunications Access Method)
  - failure as cause of takeover 613
  - failure as cause of XRF takeover 591
  - forced takeover for XRF complexes that use MNPS 601
  - interface considerations, attention notice 139
  - IRLM naming suggestions 137
  - logical unit definitions 139
  - logon message processing 601
  - master terminals, defining 647
  - mode table entry 139
  - NCP delay 141
  - ownership affecting terminal switching in XRF complex 642
  - parallel session support 139
  - parameters
    - COMCYCL 141
    - DELAY 141
    - LOGMODE 139
    - MODE 139
    - MODETBL 139
  - performance considerations 423
  - role in XRF 601
  - USERVAR variable 656
  - VARY command 139
  - XRF planning considerations 642
  - XRF process, contribution to 601
  - XRF requirements 595
- VTAM application name
  - communicating new name 606
  - defining to IMS 647
  - processing of 601, 616
  - XRF (Extended Recovery Facility) 647

- VTAM commands
  - LOGON APPLID 601
  - MODIFY 606
  - MODIFY USERVAR 616
- VTAM generic resources
  - restrictions on using 289
- VTAM terminals, encryption 341

## W

- WADS (write-ahead data set)
  - allocating 95
  - closing OLDS 95
  - definition 92
  - environment 95
  - requirement for XRF 588
- wait-for-input (WFI)
  - exclusion from program summary report 744
  - program I/O report
    - with input available 797
    - with no input available 733
  - time between transactions 829
- WLM Change State Performance Block transaction
  - current state 374
- work qualifier, definition 373
- workload
  - effect of takeover on alternate XRF system 632
  - monitoring 369, 391
  - planning for alternate XRF system 632
- workload management
  - migrating to 374
  - operating mode
    - compatibility 374
    - goal 374
- workload manager 369
  - distribution 369
  - performance objectives 370
- Workload Manager
  - business importance 372
  - performance goals 372
  - performance objectives 372
- write-ahead data set (WADS)
  - allocating 95
  - closing OLDS 95
  - definition 92
  - environment 95

## X

- X.25 terminal requirements for XRF support 595
- XRF
  - generic resources
    - MNPS 292
    - USERVAR 292
  - system failure
    - resume service 585
- XRF (Extended Recovery Facility)
  - active IMS
    - definition 586
    - failure 586
  - alternate IMS system 586

- XRF (Extended Recovery Facility)
  - (continued)
    - backup sessions 591
    - application programs 595
    - APPLID= keyword 647
    - ARM considerations 644
    - backup sessions 591
    - BACKUP= keyword 636
    - benefits 586
    - class-1 terminals
      - description 634
      - ownership of 642
    - class-2 terminals 635
    - class-3 terminals 636
    - comparison of MNPS and USERVAR 590
    - complex
      - description of 588
      - initialization phase 608
      - synchronization phase 610
      - takeover, after 591, 621
      - takeover, before 591
      - takeover, during 591, 613
      - termination, at 623
    - complex with one CPC 624
    - component roles 597
    - concepts 586
    - contributions
      - IMS 597
    - customizing individual terminals 648
    - data set placement 656
    - DB/DC environment 10
    - DBCTL capabilities 591
    - defining IMS to VTAM 656
    - defining passwords to IMS 647
    - defining VTAM application name to IMS 647
    - DFSHSBxx parameters 649
    - DFSMS contribution to XRF
      - process 601
    - DSE 586
    - DSE (dependent service element)
      - definition 597
    - dumping activity after takeover 591, 621
    - EEQE in post-takeover phase 621
    - example with one complex, two CPCs 626
    - example with one complex, two CPCs, and non-XRF IMS 626
    - example with two complexes, four CPCs 629
    - example with two complexes, three CPCs 628
    - forced takeover, VTAM 601
    - global resource serialization 644
    - hardware requirements 595
    - HSBID parameter 646
    - HSBMBR parameter 646
    - IMS master terminals 647
    - IMS secondary terminals 647
    - IMSplex 585
      - online change restriction without RM 595
    - in an IMSplex 23
    - introduction 585
    - ISC link 633

- XRF (Extended Recovery Facility)
  - (continued)
    - JES considerations 644
    - licensed program requirements 595
    - limitations 632
    - limitations of 591
    - LNK parameter for surveillance 649
    - local queue manager data sets 209
    - log close and switch 614
    - LOG parameter for surveillance 649
    - log protection 614
    - logging on 597
    - MNPS and USERVAR
      - comparison 590
    - MNPS parameter 646
    - MNPS= keyword 649
    - MNPSPW parameter 646
    - MNPSPW= keyword 649
    - MSC links 649
    - NCP
      - backup sessions 591
    - NCP contribution 605
    - NCP planning 643
    - network changes, initiating 616
    - NO parameter for surveillance 649
    - online change 443, 448
      - IMSplex restriction without RM 595
    - operational requirements 595
    - organization 624
    - overview 585
    - phases of processing 608, 624
    - planning 632, 643
    - post-takeover phase 621
    - problem determination after takeover 591, 621
    - RACF considerations 644
    - RDS parameter for surveillance 649
    - recommendation for DBRC 611
    - recoverable service element (RSE) 586, 649
    - recovery 114
    - requirements 588
      - hardware 595
      - software 595
    - returning failed active IMS as new active IMS 621
    - RSENAME= keyword 649
    - RSR and 681
    - RSR, comparison with 680
    - session recovery 648
    - signals from active IMS
      - IMS system log 597
      - ISC link 597
      - RDS 597
    - SNAPQ 75
    - software requirements 585
    - specifying backup option 636
    - specifying IMS.PROCLIB members 649
    - SSP contribution 605
    - surveillance
      - change 597
      - establishing 606
      - options 649
      - starting 597
      - stopping 597



XRF (Extended Recovery Facility)  
 (continued)  
 SWITCH= keyword for XRF 649  
 sysplex  
   global online change 449  
 system definition 645  
 system definition macros for  
   XRF 646  
 takeover  
   causes 586  
   criteria for 606  
   definition 586  
   executing new transactions 617  
   factors affecting rate of 616  
   problem determination after 591,  
     621  
   user perspective 638  
 takeover conditions 649  
 tasks performed on former active  
 system  
   code maintenance 620  
   hardware configuration  
     changes 620  
   performing hardware  
     maintenance 620  
   performing library  
     maintenance 620  
   preventive maintenance 620  
   testing backup IMS  
     procedures 620  
 terminals in 633  
 termination phase 623  
 terminology 586  
 third system as an alternate 621  
 tracking phase  
   application program and  
     transaction status 611  
   communication network  
     status 611  
   database status 611  
   dependent region status 611  
   message queue status 611  
   MFS pool status 611  
 USERVAR parameter 646  
 USERVAR table, VTAM  
   at takeover 601  
   automated by Tivoli NetView for  
     z/OS 616  
   defining IMS to VTAM 656  
 USERVAR variable  
   updating entries 606  
 USERVAR= keyword 649  
 VTAM  
   backup sessions 591  
   forced takeover 601  
 VTAM contribution to XRF  
   process 601  
 VTAM ownership affecting terminal  
   switching 642  
 z/OS contributions 600

z/OS (continued)  
 APF authorization (continued)  
   JCL considerations 125, 136  
   JES configuration 126  
   rules for 136  
 APPC/MVS administration dialog  
   updates 137  
 Automatic Restart Manager  
   (ARM) 79  
 binding 129  
 channel-to-channel (CTC) channel-end  
   appendage 132  
 commands  
   CANCEL 632  
   START 608  
 component trace (CTRACE)  
   service 380  
 components in an IMSplex 21  
 DBRC Type 4 SVC 136  
 defining IMS SVCs 130  
 failure  
   as cause of XRF takeover 591  
 IMS SVC modules 130, 131  
 installing z/OS PPT entries 126  
 interface considerations, attention  
   notice 125  
 interface modules 129  
 IRLM PPT 137  
 IRLM subsystem names  
   creating 137  
 log data set 98  
 log stream 98  
   defining 217  
 nonstandard macros 126  
 offline dump formatting 136  
 planning considerations 644  
 preventing installation problems 125  
 required IMS links to 129  
 requirement for XRF 595  
 resource cleanup module 132  
 SMF DDCONS parameter 137  
 starting availability manager 608  
 steps required to run under 129  
 system services in an IMSplex 21  
 tailoring 248  
 upgrading 202  
 VTAM Generic Resource affinity 290  
 XRF planning considerations 644  
 XRF process, contribution to 600  
 z/OS cross-system coupling facility  
   used in an IMSplex 21  
 z/OS formula in sizing structures 280  
 z/OS macro  
   DEQ 538  
   GRS 538  
   OBTAIN 538  
   RESERVE 538

## Z

z/OS  
 abend formatting 134  
 APF authorization  
   IRLM considerations 137









Product Number: 5635-A02

Printed in USA

SC19-2443-05



Spine information:

IMS    Version 11

System Administration

