



# **Program Directory for IBM Open Enterprise SDK for Apache Kafka**

V1.1.0

Program Number 5655-KAF

FMIDs HAMP110

for Use with  
z/OS

Document Date: June, 2024

GI13-5677-00

**Note**

Before using this information and the product it supports, be sure to read the general information under 7.0, “Notices” on page 25.

---

# Contents

<b>1.0 Introduction</b>	1
1.1 Open Enterprise SDK for Apache Kafka Description	1
1.2 Open Enterprise SDK for Apache Kafka FMID	2
<b>2.0 Program Materials</b>	3
2.1 Basic Machine-Readable Material	3
2.2 Optional Machine-Readable Material	4
2.3 Program Publications	4
2.3.1 Optional Program Publications	4
2.4 Program Source Materials	4
2.5 Publications Useful During Installation	4
<b>3.0 Program Support</b>	6
3.1 Program Services	6
3.2 Statement of Support Procedures	6
<b>4.0 Program and Service Level Information</b>	7
4.1 Program Level Information	7
4.2 Service Level Information	7
<b>5.0 Installation Requirements and Considerations</b>	8
5.1 Driving System Requirements	8
5.1.1 Machine Requirements	8
5.1.2 Programming Requirements	8
5.2 Target System Requirements	9
5.2.1 Machine Requirements	9
5.2.2 Programming Requirements	9
5.2.2.1 Installation Requisites	10
5.2.2.2 Operational Requisites	10
5.2.2.3 Toleration/Coexistence Requisites	11
5.2.2.4 Incompatibility (Negative) Requisites	11
5.2.3 DASD Storage Requirements	11
5.3 FMIDs Deleted	14
5.4 Special Considerations	14
<b>6.0 Installation Instructions</b>	15
6.1 Installing Open Enterprise SDK for Apache Kafka	15
6.1.1 SMP/E Considerations for Installing Open Enterprise SDK for Apache Kafka	15
6.1.2 SMP/E Options Subentry Values	15
6.1.3 SMP/E CALLLIBS Processing	16
6.1.4 Sample Jobs	16
6.1.5 Allocate and Initialize the SMP/E CSI (Optional)	18

6.1.6	Perform SMP/E RECEIVE . . . . .	18
6.1.7	Allocate SMP/E Target and Distribution Libraries . . . . .	18
6.1.8	Allocate, create and mount zFS Files (Optional) . . . . .	18
6.1.9	Allocate File System Paths . . . . .	19
6.1.10	Create DDDEF Entries . . . . .	20
6.1.11	Perform SMP/E APPLY . . . . .	20
6.1.12	Run the Installation Verification Programs . . . . .	22
6.1.13	Perform SMP/E ACCEPT . . . . .	23
<b>7.0</b>	<b>Notices . . . . .</b>	<b>25</b>
7.1	Trademarks . . . . .	25
	<b>Reader's Comments . . . . .</b>	<b>27</b>

---

## Figures

1.	Program File Content . . . . .	3
2.	Publications Useful During Installation . . . . .	4
3.	Component IDs . . . . .	6
4.	Driving System Software Requirements . . . . .	9
5.	Target System Mandatory Installation Requisites . . . . .	10
6.	Target System Mandatory Operational Requisites . . . . .	10
7.	Target System Conditional Operational Requisites . . . . .	11
8.	Total DASD Space Required by Open Enterprise SDK for Apache Kafka . . . . .	11
9.	Storage Requirements for Open Enterprise SDK for Apache Kafka Target Libraries . . . . .	13
10.	Open Enterprise SDK for Apache Kafka File System Paths . . . . .	13
11.	Storage Requirements for Open Enterprise SDK for Apache Kafka Distribution Libraries . . . . .	13
12.	SMP/E Options Subentry Values . . . . .	15
13.	Sample Installation Jobs . . . . .	16

---

## 1.0 Introduction

This program directory is intended for the system programmers who are responsible for program installation and maintenance. It contains information about the material and procedures associated with the installation of the IBM Open Enterprise SDK for Apache Kafka. This publication refers to IBM Open Enterprise SDK for Apache Kafka as Open Enterprise SDK for Apache Kafka.

The Program Directory contains the following sections:

- 2.0, “Program Materials” on page 3 identifies the basic program materials and documentation for Open Enterprise SDK for Apache Kafka.
- 3.0, “Program Support” on page 6 describes the IBM support available for Open Enterprise SDK for Apache Kafka.
- 4.0, “Program and Service Level Information” on page 7 lists the APARs (program level) and PTFs (service level) that have been incorporated into Open Enterprise SDK for Apache Kafka.
- 5.0, “Installation Requirements and Considerations” on page 8 identifies the resources and considerations that are required for installing and using Open Enterprise SDK for Apache Kafka.
- 6.0, “Installation Instructions” on page 15 provides detailed installation instructions for Open Enterprise SDK for Apache Kafka. It also describes the procedures for activating the functions of Open Enterprise SDK for Apache Kafka, or refers to appropriate publications.

Before installing Open Enterprise SDK for Apache Kafka, read the *CBPDO Memo To Users* and the *CBPDO Memo To Users Extension* that are supplied with this program in softcopy format and this program directory; then keep them for future reference.

Open Enterprise SDK for Apache Kafka is supplied in a Custom-Built Product Delivery Offering (CBPDO, 5751-CS3). The program directory that is provided in softcopy format on the CBPDO tape is identical to the hardcopy format if one was included with your order. All service and HOLDDATA for Open Enterprise SDK for Apache Kafka are included on the CBPDO tape.

Do not use this program directory if you install Open Enterprise SDK for Apache Kafka with a SystemPac or ServerPac. When you use one of those offerings, use the jobs and documentation supplied with the offering. The offering will point you to specific sections of this program directory as needed.

---

## 1.1 Open Enterprise SDK for Apache Kafka Description

### IBM Open Enterprise SDK for Apache Kafka

IBM Open Enterprise SDK for Apache Kafka 1.1 provides a straightforward connection to any Kafka cluster on the z/OS platform. z/OS application developers with COBOL or C/C++ applications, can consume and publish Kafka events from their z/OS native applications, directly to their Kafka broker.

These direct calls do not need to go through any other layer, such as Java, MQ or CICS, allowing application developers to implement their solution independently from other teams.

z/OS COBOL application developers can transform between z/OS native COBOL copybooks and event formats, such as JSON, that may be defined in schema registries. This allows the COBOL application programmer to easily include the copybook containing the event information into their applications.

With Open Enterprise SDK for Apache Kafka, clients can easily extend their current COBOL or C/C++ mainframe applications, to take advantage of Kafka event processing. Open Enterprise SDK for Apache Kafka strengthens the IBM Z platform by providing native Kafka API support for z/OS. COBOL and C/C++ samples are included in the product to help application programmers.

---

## 1.2 Open Enterprise SDK for Apache Kafka FMID

Open Enterprise SDK for Apache Kafka consists of the following FMID:

HAMP110

---

## 2.0 Program Materials

An IBM program is identified by a program number. The program number for Open Enterprise SDK for Apache Kafka is 5655-KAF.

Basic Machine-Readable Materials are materials that are supplied under the base license and are required for the use of the product.

The program announcement material describes the features supported by Open Enterprise SDK for Apache Kafka. Ask your IBM representative for this information if you have not already received a copy.

---

### 2.1 Basic Machine-Readable Material

The distribution medium for this program is physical media or downloadable files. This program is in SMP/E RELFILE format and is installed by using SMP/E. See 6.0, “Installation Instructions” on page 15 for more information about how to install the program.

You can find information about the physical media for the basic machine-readable materials for Open Enterprise SDK for Apache Kafka in the *CBPDO Memo To Users Extension*.

Figure 1 describes the program file content for Open Enterprise SDK for Apache Kafka. You can refer to the *CBPDO Memo To Users Extension* to see where the files reside on the tape.

#### Notes:

1. The data set attributes in this table must be used in the JCL of jobs that read the data sets. However, because the data sets are in IEBCOPY unloaded format, their actual attributes might be different.
2. If any RELFILES are identified as PDSEs, ensure that SMPTLIB data sets are allocated as PDSEs.

Figure 1 (Page 1 of 2). Program File Content				
Name	ORG	RECFM	RECL	BLK SIZE
SMPMCS	SEQ	FB	80	6400
IBM.HAMP110.F1	PDSE	FB	80	32720
IBM.HAMP110.F2	PDSE	VB	255	32760
IBM.HAMP110.F3	PDSE	FB	80	32720
IBM.HAMP110.F4	PDSE	U	0	32760
IBM.HAMP110.F5	PDSE	FB	80	32720

Figure 1 (Page 2 of 2). Program File Content

Name	O R G	R E C F M	L R E C L	BLK SIZE
IBM.HAMP110.F6	PDSE	FB	80	32720

## 2.2 Optional Machine-Readable Material

No optional machine-readable materials are provided for Open Enterprise SDK for Apache Kafka.

## 2.3 Program Publications

This section identifies the basic and optional publications for Open Enterprise SDK for Apache Kafka.

### 2.3.1 Optional Program Publications

No optional publications are provided for Open Enterprise SDK for Apache Kafka.

## 2.4 Program Source Materials

No program source materials or viewable program listings are provided for Open Enterprise SDK for Apache Kafka.

## 2.5 Publications Useful During Installation

You might want to use the publications listed in Figure 2 during the installation of Open Enterprise SDK for Apache Kafka.

Figure 2. Publications Useful During Installation

Publication Title	Form Number
IBM SMP/E for z/OS User's Guide	SA23-2277
IBM SMP/E for z/OS Commands	SA23-2275
IBM SMP/E for z/OS Reference	SA23-2276
IBM SMP/E for z/OS Messages, Codes, and Diagnosis	GA32-0883

**Note:** These publications can be found in IBM Documentation. Use a web browser with internet access to refer to: <https://www.ibm.com/docs/en/zos/2.5.0?topic=zos-smpe>

---

## 3.0 Program Support

This section describes the IBM support available for Open Enterprise SDK for Apache Kafka.

---

### 3.1 Program Services

Contact your IBM representative for specific information about available program services.

---

### 3.2 Statement of Support Procedures

Report any problems which you feel might be an error in the product materials to your IBM Support Center. You may be asked to gather and submit additional diagnostics to assist the IBM Support Center in their analysis.

Figure 3 lists the component IDs (COMPID) for Open Enterprise SDK for Apache Kafka.

<i>Figure 3. Component IDs</i>			
<b>FMID</b>	<b>COMPID</b>	<b>Component Name</b>	<b>RETAIN Release</b>
HAMP110	5655KAF00	IBM Open Enterprise SDK for Apache Kafka	110

---

## 4.0 Program and Service Level Information

This section identifies the program and relevant service levels of Open Enterprise SDK for Apache Kafka. The program level refers to the APAR fixes that have been incorporated into the program. The service level refers to the PTFs that have been incorporated into the program.

---

### 4.1 Program Level Information

The following APAR fixes against previous releases of Open Enterprise SDK for Apache Kafka have been incorporated into this release. They are listed by FMID:

- (N/A - this is the first release of this product)

---

### 4.2 Service Level Information

No PTFs against this release of Open Enterprise SDK for Apache Kafka have been incorporated into the product package.

You can also receive the latest HOLDDATA, and then add the **FIXCAT(IBM.PRODUCTINSTALL-REQUIREDSERVICE)** operand on your APPLY CHECK command. This will allow you to review the recommended and critical service that should be installed with your FMIDs.

---

## 5.0 Installation Requirements and Considerations

The following sections identify the system requirements for installing and activating Open Enterprise SDK for Apache Kafka. The following terminology is used:

- *Driving system*: the system on which SMP/E is executed to install the program.

The program might have specific operating system or product level requirements for using processes, such as binder or assembly utilities during the installation.

- *Target system*: the system on which the program is configured and run.

The program might have specific product level requirements, such as needing access to the library of another product for link-edits. These requirements, either mandatory or optional, might directly affect the element during the installation or in its basic or enhanced operation.

In many cases, you can use a system as both a driving system and a target system. However, you can make a separate IPL-able clone of the running system to use as a target system. The clone must include copies of all system libraries that SMP/E updates, copies of the SMP/E CSI data sets that describe the system libraries, and your PARMLIB and PROCLIB.

Use separate driving and target systems in the following situations:

- When you install a new level of a product that is already installed, the new level of the product will replace the old one. By installing the new level onto a separate target system, you can test the new level and keep the old one in production at the same time.
- When you install a product that shares libraries or load modules with other products, the installation can disrupt the other products. By installing the product onto a separate target system, you can assess these impacts without disrupting your production system.

---

### 5.1 Driving System Requirements

This section describes the environment of the driving system required to install Open Enterprise SDK for Apache Kafka.

#### 5.1.1 Machine Requirements

The driving system can run in any hardware environment that supports the required software.

#### 5.1.2 Programming Requirements

Figure 4. Driving System Software Requirements

Program Number	Product Name	Minimum VRM	Minimum Service Level will satisfy these APARs	Included in the shipped product?
5650-ZOS	z/OS	2.5.0	N/A	No
5655-ZOS	z/OS	3.1.0 or later	N/A	No

**Note:** SMP/E is a requirement for Installation and is an element of z/OS.

**Note:** Installation might require migration to later z/OS releases to be service supported. See <https://www.ibm.com/support/lifecycle/>

Open Enterprise SDK for Apache Kafka is installed into a file system, either HFS or zFS. Before installing Open Enterprise SDK for Apache Kafka, you must ensure that the target system file system data sets are available for processing on the driving system. OMVS must be active on the driving system and the target system file data sets must be mounted on the driving system.

If you plan to install Open Enterprise SDK for Apache Kafka in a zFS file system, this procedure/process requires that zFS be active on the driving system. Information on activating and using zFS can be found in z/OS Distributed File Service zSeries File System Administration, SC24-5989.

---

## 5.2 Target System Requirements

This section describes the environment of the target system required to install and use Open Enterprise SDK for Apache Kafka.

Open Enterprise SDK for Apache Kafka installs in the z/OS (Z038) SREL.

### 5.2.1 Machine Requirements

Open Enterprise SDK for Apache Kafka V1.1.0 will run on the following IBM Z servers:

- IBM z16™ (all models)
- IBM z15® (all models)
- IBM z14® (all models)
- IBM z13® (all models)

### 5.2.2 Programming Requirements

**5.2.2.1 Installation Requisites:** Installation requisites identify products that are required and *must* be present on the system or products that are not required but *should* be present on the system for the successful installation of this product.

Mandatory installation requisites identify products that are required on the system for the successful installation of this product.

Figure 5. Target System Mandatory Installation Requisites

Program Number	Product Name	Minimum VRM	Minimum Service Level will satisfy these APARs	Included in the shipped product?
5650-ZOS	z/OS	2.5.0	N/A	No
5655-ZOS	z/OS	3.1.0 or later	N/A	No

**Note:** Installation might require migration to new z/OS releases to be service supported. See [http://www.ibm.com/systems/z/os/zos/support/zos\\_eos\\_dates.html](http://www.ibm.com/systems/z/os/zos/support/zos_eos_dates.html).

Conditional installation requisites identify products that are *not* required for successful installation of this product but can resolve such things as certain warning messages at installation time.

Open Enterprise SDK for Apache Kafka has no conditional installation requisites.

**5.2.2.2 Operational Requisites:** Operational requisites are products that are required and *must* be present on the system or products that are not required but *should* be present on the system for this product to operate all or part of its functions.

Mandatory operational requisites identify products that are required for this product to operate its basic functions.

Figure 6. Target System Mandatory Operational Requisites

Program Number	Product Name and Minimum VRM/Service Level	Function
5650-ZOS	z/OS 2.5.0	
5650-ZOS	z/OS UNIX System Service for z/OS 2.5.0	
5655-ZOS	z/OS 3.1.0 or later	
5655-ZOS	z/OS UNIX System Service for z/OS 3.1.0 or later	

Conditional operational requisites identify products that are *not* required for this product to operate its basic functions but are required at run time for this product to operate specific functions.

<i>Figure 7. Target System Conditional Operational Requisites</i>		
<b>Program Number</b>	<b>Product Name and Minimum VRM/Service Level</b>	<b>Function</b>
5655-DGH	IBM Java 8 for z/OS or later	For running data transformation utility
5650-ZOS	IBM Open XL C/C++ 1.1 for z/OS with APAR(PTF) PH50622(UI83129)	For building C/C++ application
5655-EC6	IBM Enterprise COBOL for z/OS 6.4 with APAR(PTFs) PH58893(UI95101,UI95102,UI95103,UI95104)	For building COBOL application

**5.2.2.3 Toleration/Coexistence Requisites:** Toleration/coexistence requisites identify products that must be present on sharing systems. These systems can be other systems in a multisystem environment (not necessarily sysplex), a shared DASD environment (such as test and production), or systems that reuse the same DASD environment at different time intervals.

Open Enterprise SDK for Apache Kafka has no toleration or coexistence requisites.

**5.2.2.4 Incompatibility (Negative) Requisites:** Negative requisites identify products that must not be installed on the same system as this product.

Open Enterprise SDK for Apache Kafka has no negative requisites.

## 5.2.3 DASD Storage Requirements

Open Enterprise SDK for Apache Kafka libraries can reside on all supported DASD types.

Figure 8 lists the total space that is required for each type of library.

<i>Figure 8. Total DASD Space Required by Open Enterprise SDK for Apache Kafka</i>	
<b>Library Type</b>	<b>Total Space Required in 3390 Trks</b>
Target	360
Distribution	3400
File System(s)	6800

### Notes:

1. For non-RECFM U data sets, IBM recommends using system-determined block sizes for efficient DASD utilization. For RECFM U data sets, IBM recommends using a block size of 32760, which is most efficient from the performance and DASD utilization perspective.
2. Abbreviations used for data set types are shown as follows.

- U** Unique data set, allocated by this product and used by only this product. This table provides all the required information to determine the correct storage for this data set. You do not need to refer to other tables or program directories for the data set size.
- S** Shared data set, allocated by this product and used by this product and other products. To determine the correct storage needed for this data set, add the storage size given in this table to those given in other tables (perhaps in other program directories). If the data set already exists, it must have enough free space to accommodate the storage size given in this table.
- E** Existing shared data set, used by this product and other products. This data set is *not* allocated by this product. To determine the correct storage for this data set, add the storage size given in this table to those given in other tables (perhaps in other program directories). If the data set already exists, it must have enough free space to accommodate the storage size given in this table.

If you currently have a previous release of this product installed in these libraries, the installation of this release will delete the old release and reclaim the space that was used by the old release and any service that had been installed. You can determine whether these libraries have enough space by deleting the old release with a dummy function, compressing the libraries, and comparing the space requirements with the free space in the libraries.

For more information about the names and sizes of the required data sets, see 6.1.7, “Allocate SMP/E Target and Distribution Libraries” on page 18.

3. Abbreviations used for the file system path type are as follows.

- N** New path, created by this product.
- X** Path created by this product, but might already exist from a previous release.
- P** Previously existing path, created by another product.

4. All target and distribution libraries listed have the following attributes:

- The default name of the data set can be changed.
- The default block size of the data set can be changed.
- The data set can be merged with another data set that has equivalent characteristics.
- The data set can be either a PDS or a PDSE.

5. All target libraries listed have the following attributes:

- These data sets can be SMS-managed, but they are not required to be SMS-managed.
- These data sets are not required to reside on the IPL volume.
- The values in the "Member Type" column are not necessarily the actual SMP/E element types that are identified in the SMPMCS.

6. All target libraries that are listed and contain load modules have the following attributes:

- These data sets can be in the LPA, but they are not required to be in the LPA.
- These data sets can be in the LNKLIST.
- These data sets are not required to be APF-authorized.
- Open Enterprise SDK for Apache Kafka requires that the SMPLTS data set must be a PDSE. If your existing SMPLTS is a PDS, you will need to allocate a new PDSE and copy your existing SMPLTS into it and then change the SMPLTS DDDEF entry to indicate the new PDSE data set.

7. The total space required listed in Figure 8 on page 11 for File System is accounted sufficiently for Open Enterprise SDK for Apache Kafka at base release. On the other hand, due to the natural growth in size for Open Enterprise SDK for Apache Kafka during PTFs to avoid encountering insufficient space quota problems, it is advised to allocate File Systems that can be auto-extended.

The following figures describe the target and distribution libraries and file system paths required to install Open Enterprise SDK for Apache Kafka. The storage requirements of Open Enterprise SDK for Apache Kafka must be added to the storage required by other programs that have data in the same library or path.

**Note:** Use the data in these tables to determine which libraries can be merged into common data sets. In addition, since some ALIAS names may not be unique, ensure that no naming conflicts will be introduced before merging libraries.

Figure 9. Storage Requirements for Open Enterprise SDK for Apache Kafka Target Libraries

Library DDNAME	Member Type	Target Volume	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
SIXYCOPY	SAMP	T2	U	LIBRARY	FB	80	15	2
SIXYJCL	SAMP	T2	U	LIBRARY	FB	80	15	2
SIXYLOAD	PROGRAM	T2	U	LIBRARY	U	0	15000	4
SIXYSAMP	SAMP	T2	U	LIBRARY	FB	80	15	2
SIXYSDEF	DATA	T2	U	LIBRARY	FB	80	500	20

Figure 10. Open Enterprise SDK for Apache Kafka File System Paths

DDNAME	T Y P E	Path Name
SIXYZ11	N	/usr/lpp/IBM/ixy/v1r1m0/IBM

Figure 11 (Page 1 of 2). Storage Requirements for Open Enterprise SDK for Apache Kafka Distribution Libraries

Library DDNAME	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
AIXYCOPY	U	LIBRARY	FB	80	15	2
AIXYJCL	U	LIBRARY	FB	80	15	2
AIXYLOAD	U	LIBRARY	U	0	15000	4

Figure 11 (Page 2 of 2). Storage Requirements for Open Enterprise SDK for Apache Kafka Distribution Libraries

Library DDNAME	T Y P E	O R G	R E C F M	L R E C L	No. of 3390 Trks	No. of DIR Blks
AIXYSAMP	U	LIBRARY	FB	80	15	2
AIXYSDEF	U	LIBRARY	FB	80	500	20
AIXYSRC	U	LIBRARY	VB	255	500	2

## 5.3 FMIDs Deleted

Installing Open Enterprise SDK for Apache Kafka might result in the deletion of other FMIDs. To see which FMIDs will be deleted, examine the ++VER statement in the SMPMCS of the product.

If you do not want to delete these FMIDs at this time, install Open Enterprise SDK for Apache Kafka into separate SMP/E target and distribution zones.

**Note:** These FMIDs are not automatically deleted from the Global Zone. If you want to delete these FMIDs from the Global Zone, use the SMP/E REJECT NOFMID DELETEFMID command. See the SMP/E Commands book for details.

## 5.4 Special Considerations

There are no special considerations for this product.

---

## 6.0 Installation Instructions

This chapter describes the installation method and the step-by-step procedures to install and to activate the functions of Open Enterprise SDK for Apache Kafka.

Notes:

- You can use the sample jobs that are provided to perform part or all of the installation tasks. The SMP/E jobs assume that all DDDEF entries that are required for SMP/E execution have been defined in appropriate zones.
- You can use the SMP/E dialogs instead of the sample jobs to accomplish the SMP/E installation steps.

---

### 6.1 Installing Open Enterprise SDK for Apache Kafka

#### 6.1.1 SMP/E Considerations for Installing Open Enterprise SDK for Apache Kafka

Use the SMP/E RECEIVE, APPLY, and ACCEPT commands to install this release of Open Enterprise SDK for Apache Kafka.

Note: The default High Level Qualifier prefix used in the install JCL for SMP/E datasets, distribution libraries, and target libraries is CBC.HAMP110, change this if you have chosen a different SMPPRFX, DSTPRFX, or TGTPRFX respectively. The recommended default high level directory name for zFS path is root "/", update #PathPrefix in the respective install JCL sample respectively or change this if you have chosen a different directory path. If a non-default prefix is chosen for the #PathPrefix for zFS directory, make sure to prepend the paths in the IVT steps in 6.1.12 with the chosen #PathPrefix

#### 6.1.2 SMP/E Options Subentry Values

The recommended values for certain SMP/E CSI subentries are shown in Figure 12. Using values lower than the recommended values can result in failures in the installation. DSSPACE is a subentry in the GLOBAL options entry. PEMAX is a subentry of the GENERAL entry in the GLOBAL options entry. See the SMP/E manuals for instructions on updating the global zone.

<i>Figure 12. SMP/E Options Subentry Values</i>		
Subentry	Value	Comment
DSSPACE	(4500,500,500)	3390 DASD tracks
PEMAX	SMP/E Default	IBM suggests using the SMP/E default for PEMAX.

### 6.1.3 SMP/E CALLLIBS Processing

Open Enterprise SDK for Apache Kafka uses the CALLLIBS function provided in SMP/E to resolve external references during installation.

### 6.1.4 Sample Jobs

The following sample installation jobs are provided as part of the product to help you install Open Enterprise SDK for Apache Kafka:

Figure 13. Sample Installation Jobs			
Job Name	Job Type	Description	RELFILE
IXYWSMPE	SMP/E	Sample job to define and prime a new SMP/E CSI ( <b>Optional</b> )	IBM.HAMP110.F1
IXYWRECV	RECEIVE	Sample RECEIVE job	IBM.HAMP110.F1
IXYWALOC	ALLOCATE	Sample job to allocate target and distribution libraries	IBM.HAMP110.F1
IXYWZFS	ALLOMZFS	Sample job to allocate, create mountpoint, & mount zFS data sets ( <b>Optional</b> )	IBM.HAMP110.F1
IXYISMKD	MKDIR	Sample job to invoke the supplied IXYMKDIR EXEC to allocate file system paths	IBM.HAMP110.F1
IXYWDDEF	DDDEF	Sample job to define SMP/E DDDEFs	IBM.HAMP110.F1
IXYWAPLY	APPLY	Sample APPLY job	IBM.HAMP110.F1
IXYWACPT	ACCEPT	Sample ACCEPT job	IBM.HAMP110.F1

You can access the sample installation jobs by performing an SMP/E RECEIVE (refer to 6.1.6, “Perform SMP/E RECEIVE” on page 18) then copy the jobs from the RELFILES to a work data set for editing and submission. See Figure 13 to find the appropriate relfile data set.

You can also copy the sample installation jobs from the tape or product files by submitting the following job. Depending on your distribution medium, use either the //TAPEIN or the //FILEIN DD statement and comment out or delete the other statement. Before you submit the job, add a job card and change the lowercase parameters to uppercase values to meet the requirements of your site.

```
//STEP1    EXEC PGM=IEBCOPY
//SYSPRINT DD SYSOUT=*
//*****
/* Make the //TAPEIN DD statement below active if you install*
/* from a CBPDO tape by uncommenting the DD statement below. *
//*****
/*TAPEIN   DD DSN=IBM.HAMP110.F1,UNIT=tunit,
/*          VOL=SER=volser,LABEL=(x,SL),
/*          DISP=(OLD,KEEP)
```

```

//*****
/* Make the //TAPEIN DD statement below active if you install*
/* from a product tape received outside the CBPDO process   *
/* (using the optional SMP/E RECEIVE job) by uncommenting   *
/* the DD statement below.                                   *
//*****
/*TAPEIN   DD DSN=IBM.HAMP110.F1,UNIT=tunit,
/*          VOL=SER=AMP110,LABEL=(2,SL),
/*          DISP=(OLD,KEEP)
//*****
/* Make the //FILEIN DD statement below active for          *
/* downloaded DASD files.                                    *
//*****
/*FILEIN   DD DSN=IBM.HAMP110.F1,UNIT=SYSALLDA,DISP=SHR,
/*          VOL=SER=filevol
//OUT      DD DSNAME=jcl-library-name,
//          DISP=(NEW,CATLG,DELETE),
//          VOL=SER=dasdvol,UNIT=SYSALLDA,
//          SPACE=(TRK,(20,10,5))
//SYSUT3   DD UNIT=SYSALLDA,SPACE=(CYL,(1,1))
//SYSIN    DD *
            COPY INDD=xxxxIN,OUTDD=OUT
            SELECT MEMBER=(IXYISMKD,IXYWACPT,IXYWALOC,IXYWAPLY)
            SELECT MEMBER=(IXYWDEF)
            SELECT MEMBER=(IXYWRECV,IXYWSMPE,IXYWZFS)
/*

```

See the following information to update the statements in the previous sample:

**TAPEIN:**

**tunit** is the unit value that matches the product package.

**volser** is the volume serial that matches the product package.

**x** is the tape file number that indicates the location of the data set name on the tape.

See the documentation that is provided by CBPDO for the location of IBM.HAMP110.F1 on the tape.

**FILEIN:**

**filevol** is the volume serial of the DASD device where the downloaded files reside.

**OUT:**

**jcl-library-name** is the name of the output data set where the sample jobs are stored.

**dasdvol** is the volume serial of the DASD device where the output data set resides.

**SYSIN:**

**xxxxIN** is either TAPEIN or FILEIN depending on your input DD statement.

### 6.1.5 Allocate and Initialize the SMP/E CSI (Optional)

If you are using an existing CSI do not execute this job.

If you install into existing SMP/E data sets, make sure that you have enough space.

If you plan to install into an existing zone, the cluster should already have been allocated and primed. You can go on to the next step to perform a SMP/E RECEIVE.

To install into a new zone, use the IXYWSMPE sample job to allocate and prime the SMPCSI cluster. Consult the instructions in the sample job for more information.

**Expected Return Codes and Messages:** You will receive a return code of 0 if this job runs correctly.

### 6.1.6 Perform SMP/E RECEIVE

If you have obtained Open Enterprise SDK for Apache Kafka as part of a CBPDO, use the RCVPDO job in the CBPDO RIMLIB data set to receive the Open Enterprise SDK for Apache Kafka FMIDs, service, and HOLDDATA that are included on the CBPDO package. For more information, see the documentation that is included in the CBPDO.

You can also choose to edit and submit sample job IXYWRECV to perform the SMP/E RECEIVE for Open Enterprise SDK for Apache Kafka. Consult the instructions in the sample job for more information.

Note: By default the RECEIVE job will look for the relfiles under IBM.HAMP110.\* If instead they are stored under #relhlq.IBM.HAMP110.\*, RFPREFIX(#relhlq) will need to be added to the RECEIVE job. For example: RECEIVE S( HAMP110 ) RFPREFIX(#relhlq) SYSMODS LIST.

**Expected Return Codes and Messages:** You will receive a return code of 0 if this job runs correctly.

### 6.1.7 Allocate SMP/E Target and Distribution Libraries

Edit and submit sample job IXYWALOC to allocate the SMP/E target and distribution libraries for Open Enterprise SDK for Apache Kafka. Consult the instructions in the sample job for more information.

**Expected Return Codes and Messages:** You will receive a return code of 0 if this job runs correctly.

### 6.1.8 Allocate, create and mount zFS Files (Optional)

This job allocates, creates a mountpoint, and mounts zFS data sets.

If you plan to install Open Enterprise SDK for Apache Kafka into a new z/OS UNIX file system, you can edit and submit the optional IXYWZFS job to perform the following tasks:

- Create the z/OS UNIX file system

- Create a mount point
- Mount the z/OS UNIX file system on the mountpoint

Consult the instructions in the sample job for more information.

The recommended z/OS UNIX file system type is *zFS*. The recommended mount point is */usr/lpp/IBM/ixy/v1r1m0*.

Before running the sample job to create the z/OS UNIX file system, you must ensure that OMVS is active on the driving system. *zFS* must be active on the driving system if you are installing Open Enterprise SDK for Apache Kafka into a file system that is *zFS*.

If you create a new file system for this product, consider updating the BPXPRMxx PARMLIB member to mount the new file system at IPL time. This action can be helpful if an IPL occurs before the installation is completed.

```
MOUNT FILESYSTEM('#dsn')
MOUNTPOINT('/usr/lpp/IBM/ixy/v1r1m0')
MODE(RDWR) /* can be MODE(READ) */
TYPE(ZFS) PARM('AGGRGROW') /* zFS, with extents */
```

See the following information to update the statements in the previous sample:

**#dsn** is the name of the data set holding the z/OS UNIX file system.

**/usr/lpp/IBM/ixy/v1r1m0** is the name of the mount point where the z/OS UNIX file system will be mounted.

**Expected Return Codes and Messages:** You will receive a return code of 0 if this job runs correctly.

## 6.1.9 Allocate File System Paths

The target system HFS or *zFS* data set must be mounted on the driving system when running the sample IXYISMKD job since the job will create paths in the HFS or *zFS*.

Before running the sample job to create the paths in the file system, you must ensure that OMVS is active on the driving system and that the target system's HFS or *zFS* file system is mounted to the driving system. *zFS* must be active on the driving system if you are installing Open Enterprise SDK for Apache Kafka into a file system that is *zFS*.

If you plan to install Open Enterprise SDK for Apache Kafka into a new HFS or *zFS* file system, you must create the mountpoint and mount the new file system to the driving system for Open Enterprise SDK for Apache Kafka.

The recommended mountpoint is */usr/lpp/IBM/ixy/v1r1m0*.

Edit and submit sample job IXYISMKD to allocate the HFS or *zFS* paths for Open Enterprise SDK for Apache Kafka. Consult the instructions in the sample job for more information.

If you create a new file system for this product, consider updating the BPXPRMxx PARMLIB member to mount the new file system at IPL time. This action can be helpful if an IPL occurs before the installation is completed.

**Expected Return Codes and Messages:** You will receive a return code of 0 if this job runs correctly.

### 6.1.10 Create DDDEF Entries

Edit and submit sample job IXYWDDEF to create DDDEF entries for the SMP/E target and distribution libraries for Open Enterprise SDK for Apache Kafka. Consult the instructions in the sample job for more information.

**Expected Return Codes and Messages:** You will receive a return code of 0 if this job runs correctly.

### 6.1.11 Perform SMP/E APPLY

1. Ensure that you have the latest HOLDDATA; then edit and submit sample job IXYWAPLY to perform an SMP/E APPLY CHECK for Open Enterprise SDK for Apache Kafka. Consult the instructions in the sample job for more information.

The latest HOLDDATA is available through several different portals, including <https://www.ibm.com/support/pages/enhanced-holddata-zos>. The latest HOLDDATA may identify HIPER and FIXCAT APARs for the FMIDs you will be installing. An APPLY CHECK will help you determine if any HIPER or FIXCAT APARs are applicable to the FMIDs you are installing. If there are any applicable HIPER or FIXCAT APARs, the APPLY CHECK will also identify fixing PTFs that will resolve the APARs, if a fixing PTF is available.

You should install the FMIDs regardless of the status of unresolved HIPER or FIXCAT APARs. However, do not deploy the software until the unresolved HIPER and FIXCAT APARs have been analyzed to determine their applicability. That is, before deploying the software either ensure fixing PTFs are applied to resolve all HIPER or FIXCAT APARs, or ensure the problems reported by all HIPER or FIXCAT APARs are not applicable to your environment.

To receive the full benefit of the SMP/E Causer SYSMOD Summary Report, do *not* bypass the PRE, ID, REQ, and IFREQ on the APPLY CHECK. The SMP/E root cause analysis identifies the cause only of *errors* and not of *warnings* (SMP/E treats bypassed PRE, ID, REQ, and IFREQ conditions as warnings, instead of errors).

Here are sample APPLY commands:

- a. To ensure that all recommended and critical service is installed with the FMIDs, receive the latest HOLDDATA and use the APPLY CHECK command as follows

```
APPLY S(fmid,fmid,...) CHECK
FORFMID(fmid,fmid,...)
SOURCEID(RSU*)
FIXCAT(IBM.ProductInstall-RequiredService)
GROUPEXTEND .
```

Some HIPER APARs might not have fixing PTFs available yet. You should analyze the symptom flags for the unresolved HIPER APARs to determine if the reported problem is applicable to your environment and if you should bypass the specific ERROR HOLDS in order to continue the installation of the FMIDs.

This method requires more initial research, but can provide resolution for all HIPERs that have fixing PTFs available and are not in a PE chain. Unresolved PEs or HIPERs might still exist and require the use of BYPASS.

- b. To install the FMIDs without regard for unresolved HIPER APARs, you can add the BYPASS(HOLDCLASS(HIPER)) operand to the APPLY CHECK command. This will allow you to install FMIDs even though one or more unresolved HIPER APARs exist. After the FMIDs are installed, use the SMP/E REPORT ERRSYSMODS command to identify unresolved HIPER APARs and any fixing PTFs.

```
APPLY S(fmid,fmid,...) CHECK
FORFMID(fmid,fmid,...)
SOURCEID(RSU*)
FIXCAT(IBM.ProductInstall-RequiredService)
GROUPEXTEND
BYPASS(HOLDCLASS(HIPER),HOLDFIXCAT) .
..any other parameters documented in the program directory
```

This method is quicker, but requires subsequent review of the Exception SYSMOD report produced by the REPORT ERRSYSMODS command to investigate any unresolved HIPERs. If you have received the latest HOLDDATA, you can also choose to use the REPORT MISSINGFIX command and specify Fix Category IBM.ProductInstall-RequiredService to investigate missing recommended service.

If you bypass HOLDS during the installation of the FMIDs because fixing PTFs are not yet available, you can be notified when the fixing PTFs are available by using the APAR Status Tracking (AST) function of ServiceLink or the APAR Tracking function of ResourceLink.

2. After you take actions that are indicated by the APPLY CHECK, remove the CHECK operand and run the job again to perform the APPLY.

**Note:** The GROUPEXTEND operand indicates that SMP/E applies all requisite SYSMODs. The requisite SYSMODS might be applicable to other functions.

**Expected Return Codes and Messages from APPLY CHECK:** You will receive a return code of 0 if this job runs correctly.

**Expected Return Codes and Messages from APPLY:** You will receive a return code of 0 if this job runs correctly.

## 6.1.12 Run the Installation Verification Programs

Follow the below steps to setup sample Producer and Consumer in COBOL. The following JCL samples can be found in target dataset SIXYSAMP:

- IXYJCNTL: Modify this JCL with the following changes to create the link-edit cards for the sample COBOL user defined functions - IXYCA2E, IXYCE2A and IXYDLCB:
  1. Replace @@JOB CARD@@ to a valid job card based on the environment.
  2. Change all @@IXYHLQ@@ to the HLQ for the IBM Open Enterprise SDK for Apache Kafka installation PDSEs.
  3. Change all @@UNIT@@ to unit type and @@VOL@@ to volume serial. These are optional. If not required delete these lines from the JCL.
  4. Submit the JCL IXYJCNTL.
- IXYCA2E: Modify the COBOL user defined function to convert ASCII string to EBCDIC string:
  1. Update value of variable - WS-CCSID-ASC to the CCSID of ASCII in your environment. Default value is set to 819.
  2. Update value of variable - WS-CCSID-EBC to the CCSID of EBCDIC in your environment. Default value is set to 1047.
  3. Maximum string length to be converted is set to 1024. If this needs updating modify the length of the variables - NATIONAL-DATA, ASCII-DATA and EBCDIC-DATA-31.
- IXYCE2A: Modify the COBOL user defined function to convert EBCDIC string to ASCII string:
  1. Update value of variable - WS-CCSID-ASC to the CCSID of ASCII in your environment. Default value is set to 819.
  2. Update value of variable - WS-CCSID-EBC to the CCSID of EBCDIC in your environment. Default value is set to 1047.
  3. Maximum string length to be converted is set to 1024. If this needs updating modify the length of the variables - NATIONAL-DATA, ASCII-DATA-31 and EBCDIC-DATA.
- IXYJSFNC: Modify this JCL with the following changes to compile and link-edit the functions - IXYCA2E, IXYCE2A and IXYDLCB:
  1. Replace @@JOB CARD@@ to a valid job card based on the environment.
  2. Change all @@IXYHLQ@@ to the HLQ for the IBM Open Enterprise SDK for Apache Kafka installation PDSEs.
  3. Change @@CEEHLQ@@ to the HLQ for Language Environment runtime libraries.
  4. Change @@IGYHLQ@@ to the HLQ for COBOL compiler library.
  5. Change all @@UNIT@@ to unit type and @@VOL@@ to volume serial. These are optional. If not required delete these lines from the JCL.
  6. Submit the JCL IXYJSFNC.

- **IXYJVER:** Modify this JCL with the following changes to compile and link-edit the sample installation verification program - IXYSVER.
  1. Replace @@JOB CARD@@ to a valid job card based on the environment.
  2. Change all @@IXYHLQ@@ to the HLQ for the IBM Open Enterprise SDK for Apache Kafka installation PDSEs.
  3. Change @@CEEHLQ@@ to the HLQ for Language Environment runtime libraries.
  4. Change @@IGYHLQ@@ to the HLQ for COBOL compiler library.
  5. Change all @@UNIT@@ to unit type and @@VOL@@ to volume serial. These are optional. If not required delete these lines from the JCL.
  6. Submit the JCL IXYJVER.
- **IXYJRNVR:** Modify this JCL with the following changes to execute the sample installation verification program - IXYSVER:
  1. Replace @@JOB CARD@@ to a valid job card based on the environment.
  2. Change all @@IXYHLQ@@ to the HLQ for the IBM Open Enterprise SDK for Apache Kafka installation PDSEs.
  3. Change @@CEEHLQ@@ to the HLQ for Language Environment runtime libraries.
  4. Change LIBPATH in CEEOPTS DD to the USS path of librdkafka.so.1.
  5. Submit the JCL IXYJRNVR.
  6. Check the job log, if no errors, then the Kafka Version will be printed in SYSOUT.

**Expected Return Codes and Messages:** You will receive the correct version number and the return code will be 0 if this job runs correctly.

### 6.1.13 Perform SMP/E ACCEPT

Edit and submit sample job IXYWACPT to perform an SMP/E ACCEPT CHECK for Open Enterprise SDK for Apache Kafka. Consult the instructions in the sample job for more information.

To receive the full benefit of the SMP/E Causer SYSMOD Summary Report, do *not* bypass the PRE, ID, REQ, and IFREQ on the ACCEPT CHECK. The SMP/E root cause analysis identifies the cause of *errors* but not *warnings* (SMP/E treats bypassed PRE, ID, REQ, and IFREQ conditions as warnings rather than errors).

Before you use SMP/E to load new distribution libraries, it is recommended that you set the ACCJCLIN indicator in the distribution zone. In this way, you can save the entries that are produced from JCLIN in the distribution zone whenever a SYSMOD that contains inline JCLIN is accepted. For more information about the ACCJCLIN indicator, see the description of inline JCLIN in the SMP/E Commands book for details.

After you take actions that are indicated by the ACCEPT CHECK, remove the CHECK operand and run the job again to perform the ACCEPT.

**Note:** The GROUPEXTEND operand indicates that SMP/E accepts all requisite SYSMODs. The requisite SYSMODS might be applicable to other functions.

**Expected Return Codes and Messages from ACCEPT CHECK:** You will receive a return code of 0 if this job runs correctly.

If PTFs that contain replacement modules are accepted, SMP/E ACCEPT processing will link-edit or bind the modules into the distribution libraries. During this processing, the Linkage Editor or Binder might issue messages that indicate unresolved external references, which will result in a return code of 4 during the ACCEPT phase. You can ignore these messages, because the distribution libraries are not executable and the unresolved external references do not affect the executable system libraries.

**Expected Return Codes and Messages from ACCEPT:** You will receive a return code of 0 if this job runs correctly.

---

## 7.0 Notices

This information was developed for products and services offered in the U.S.A. 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.

APAR numbers are provided in this document to assist in locating PTFs that may be required. Ongoing problem reporting may result in additional APARs being created. Therefore, the APAR lists in this document may not be complete. To obtain current service recommendations and to identify current product service requirements, always contact the IBM Customer Support Center.

IBM may have patents or pending patent applications covering subject matter 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 the

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, New York 10504-1785  
USA

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

---

## 7.1 Trademarks

IBM, the IBM logo, and [ibm.com](http://ibm.com) 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 "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

UNIX is a registered trademark of The Open Group in the United States and other countries.

Kafka® is a registered trademark of the Apache Software Foundation in the United States. No endorsement by The Apache Software Foundation is implied by using this mark.

---

## Reader's Comments

### Program Directory for IBM Open Enterprise SDK for Apache Kafka, June, 2024

We appreciate your input on this publication. Feel free to comment on the clarity, accuracy, and completeness of the information or give us any other feedback that you might have.

Send your comments by emailing us at [ibmdocs@us.ibm.com](mailto:ibmdocs@us.ibm.com), and include the following information:

Your name and address  
Your email address  
Your telephone or fax number  
The publication title and order number  
The topic and page number related to your comment  
The text of your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you submit.

Thank you for your participation.



Printed in USA

GI13-5677-00

