Platform Analytics
Version 9.1
for LSF

Installing

SC14-7571-00

Note
Before using this information and the product it supports, read the information in "Notices" on page 45.

First edition

This edition applies to version 9, release 1, modification 0 of Platform Analytics (product number 5725-G84) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Figures

Tables

Chapter 1. Understanding Platform Analytics

Get to know the components, ports, and architecture of IBM® Platform Analytics.

Platform Analytics hosts

These are the types of Platform Analytics hosts that you can install:

Analytics node

Analytics nodes are hosts on which Platform Analytics Data Collectors for LSF is installed and which collect operational data from Platform LSF clusters or license servers. Each node belongs either to a cluster from which Platform Analytics is to collect data (including license data, if connected to a license server), or is a stand-alone host that collects license data.

Analytics database

The Analytics database runs on one or more hosts and stores the cluster operational data loaded from the Analytics nodes for later reporting and analysis.

Analytics server

The Analytics server manages the data that the Analytics nodes collect. You can perform all server functions using the Analytics Console in the server.

Analytics reporting server

The Analytics reporting server generates reports based on the data that the Analytics nodes collect. You can view these reports using a web browser. The Analytics reporting server can run on the same host as the Analytics server if this host meets the system requirements of both the Analytics server and the Tableau Server software.

System port usage

Platform Analytics hosts and components use the following system ports. You must ensure that the system ports for Platform Analytics are open for the hosts and components to communicate with one another. You must also ensure that your web server port is open.

Analytics system ports

Platform Analytics hosts use the following TCP ports by default. The following tables describe the configuration files that you must modify to change the default settings and the firewall configurations you must set to allow communication using these ports between the various hosts.

Analytics server ports

Table 1 lists information about the default Analytics server ports.

Table 1. Default Platform Analytics server ports

Port name	Default port number	Description	Firewall configuration (all bidirectional)	Configuration file
PIAM_PORT	9091	Internal port for the task scheduler.	not applicable	*ANALYTICS_TOP*\conf\pi.conf
PIEM_PORT	9092	The Analytics event manager uses this port to receive events from the Analytics server and nodes.	Allow connections for this port from each node host to the server host.	*ANALYTICS_TOP*\conf\pi.conf
Remoting server port	9093	The Analytics remoting server uses this port to receive data from the remoting node.	Allow connections for this port from the remoting node to the server host.	*ANALYTICS_TOP*\conf\ remotingserver.xml

Analytics node ports

Table 2 lists information about the default Analytics node ports.

Table 2. Default Platform Analytics node ports

Port name	Default port number	Description	Firewall configuration (all bidirectional)	Configuration file
PLC port	4046	The Analytics server uses this port to manage the Platform loader controller (**plc**) on the Analytics node.	Allow connections for this port from the server host to each node host.	*ANALYTICS_TOP*/conf/plc.xml/ conf/plc.xml

Third-party tool ports

Table 3 lists information about the default Analytics third-party tool ports.

Table 3. Default Platform Analytics third-party tool ports

Port name	Default port number	Description	Firewall configuration (all bidirectional)
Web server port	80	Users can use this port to browse reports.	Allow connections for this port from each browsing host to the server host.
License server port	N/A	To collect license usage data from the FLEXnet server, the Analytics node has to communicate with the FLEXnet server host via TCP.	Allow connections for this port from the node host to each FLEXnet server host.
Vertica port	5433	This port is required for client connectivity, such as JDBC.	Allow connections for this port from server and node hosts.
FLEXnet Manager database port	1433 for MS SQL Server	To collect data from FLEXnet Manager, the Analytics node has to communicate with the FLEXnet Manager database via TCP.	Allow connections for this port from the node host to the FLEXnet Manager database host.

System architecture diagrams

Cluster operations data collected from the Analytics nodes is stored in the Analytics database for subsequent reporting and analysis by IBM Platform Analytics to provide the information you need to make business decisions, such as:

- Identify and remove resource bottlenecks that are affecting business
- Plan for capacity and future needs by seeing how capacity is being used and where capacity is available
- Optimize asset utilization by seeing how your infrastructure and resources are being used
- Improve productivity and efficiency of your applications, workload costs, and computing environment

The following diagrams show sample architectures for Platform Analytics systems that monitor multiple IBM Platform LSF clusters.

Figure 1 on page 4 illustrates a system architecture where each Analytics node is a member of an LSF® cluster.

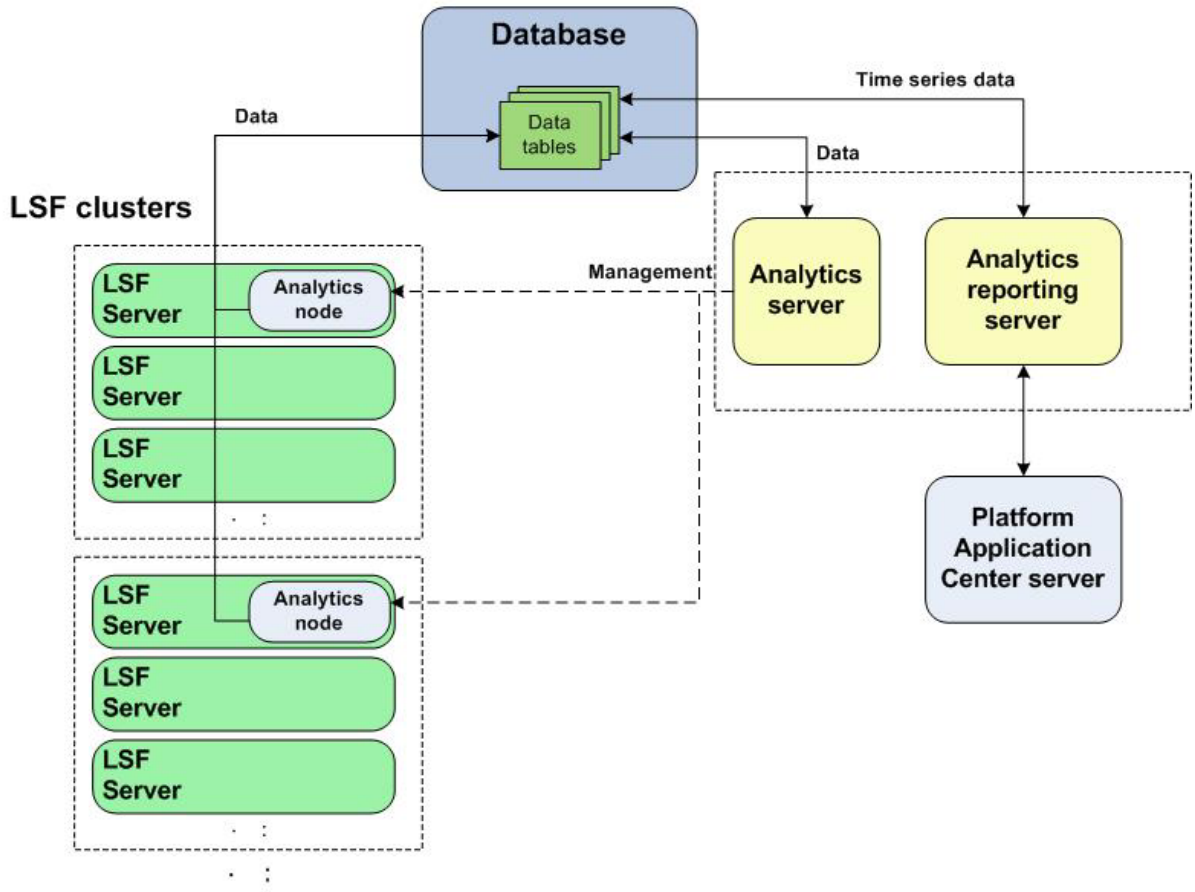

Figure 1. Platform Analytics system architecture where each Analytics node is a member of an LSF cluster

Figure 2 on page 5 illustrates a system architecture with a stand-alone node connected to a license server.

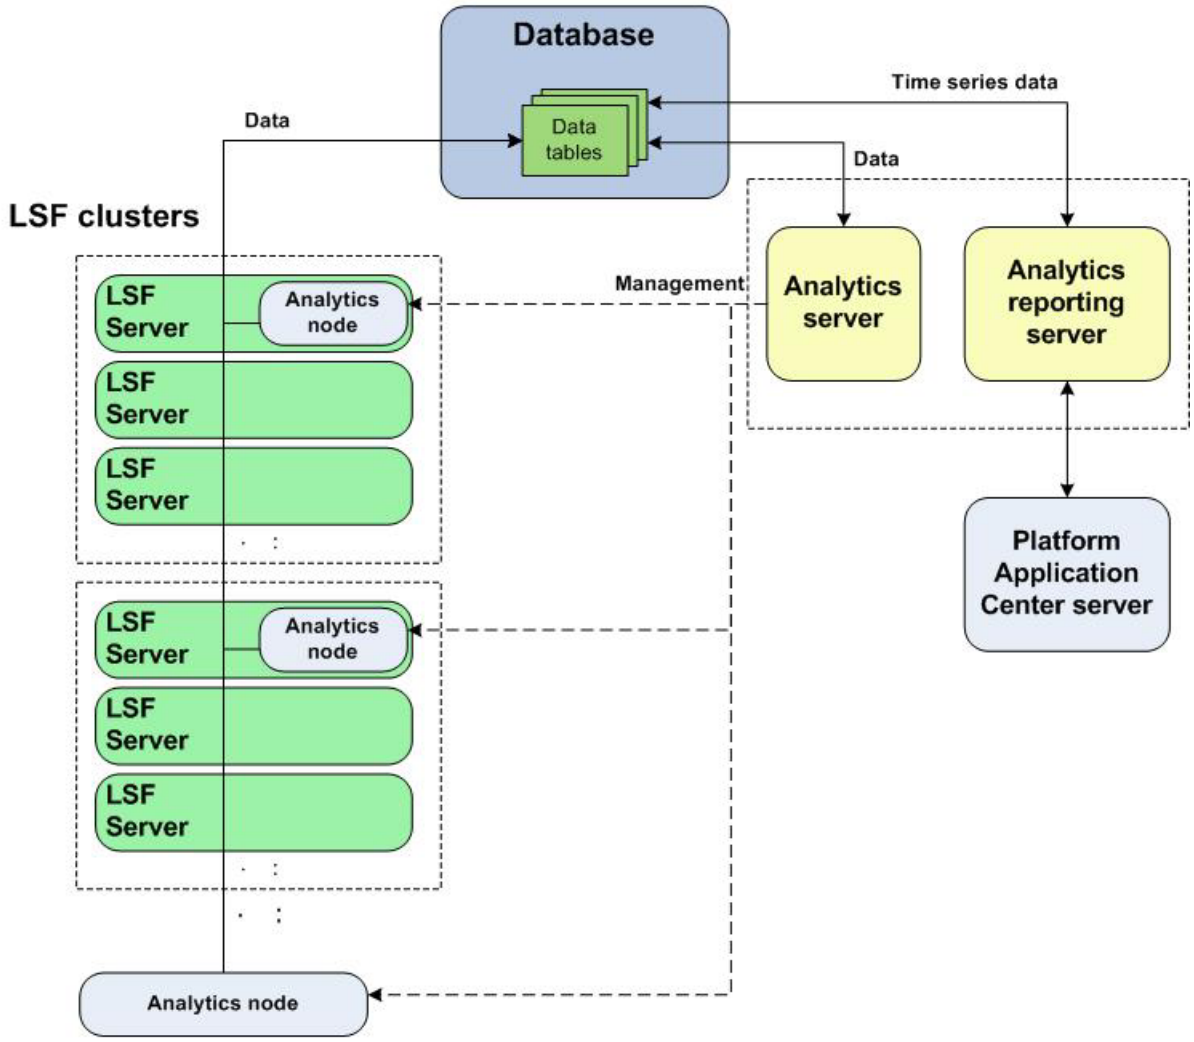

Figure 2. Platform Analytics system architecture with a stand-alone node connected to a license server

Chapter 2. Preparing to install Platform Analytics

Before installing Platform Analytics, you will select the hosts that meet the detailed system requirements, prepare the IBM Platform LSF cluster, and download the Platform Analytics installation packages.

Selecting the hosts

You must select the appropriate hosts that meet the detailed system requirements for the Analytics database hosts, Analytics server host, Platform Analytics reporting server host, and Analytics node hosts.

Selecting the database hosts

For optimal performance of your production database, the Analytics database cluster should consist of at least three dedicated multi-core hosts running on a high-bandwidth network. Since the Analytics database needs to share a large volume of data among the database nodes in the database cluster during data loading or data querying, network bandwidth is an important performance bottleneck for a production database. Therefore, the Analytics database cluster should have a Gigabit Ethernet connection with the Platform Analytics reporting server and the Analytics node hosts.

Vertica recommends a 1-10 GB full duplex switch for the private network interface and a VLAN or separate switch for the public network. The switch used for the private network should have sufficient bandwidth to enable 1 GB transfer speeds between any pair of nodes.

The hardware requirements are the same for all the intended database hosts. Refer to the Vertica Analytic Database documentation and *IBM Platform Analytics Release Notes* for the latest list of system requirements and supported operating systems for the Analytics database hosts.

Input/output (I/O) performance is important for the operation of a database while fault tolerance is important to safeguard your data. Using a RAID 01 or 10 system enables the database host to take advantage of data striping and data mirroring. Data striping allows data to be transferred to multiple hard disks concurrently, which improves I/O performance. Data mirroring means that your database does not lose data even if a hard disk fails.

Table 4 describes the optimal configuration of the database server depending on the size of your cluster. The hardware requirements are the same for each database host.

Table 4. Recommended hardware specifications for the Analytics database hosts

Cluster size	Number of hosts	RAM	CPU	Local hard disk	Network
Medium (100 - 1000 hosts)	3	16 GB	4 × 2.4GHz	10000 RPM SATA/SCSI/SAS/SSD RAID 01 or 10 300 GB	Gigabit Ethernet

Table 4. Recommended hardware specifications for the Analytics database hosts (continued)

Cluster size	Number of hosts	RAM	CPU	Local hard disk	Network
Large (more than 1000 hosts)	more than 3	32 GB	8 × 2.4GHz	10000 RPM SATA/SCSI/SAS/SSD RAID 01 or 10 1 TB	Gigabit Ethernet

Data striping

Data striping is the technique of segmenting logically-sequential data, such as a single file, so that the database can assign segments to multiple physical devices (usually disk drives for RAID storage, or network interfaces for grid-oriented storage) in a round-robin fashion and thus be read or written concurrently.

Automatic data striping is available in certain RAID devices under software or hardware control, and in file systems of clusters. The following parameters are important when improving I/O performance:

Stripe width

The number of parallel stripes that can be written to or read from simultaneously. This is the number of disks in the RAID system, and as it increases, the read/write performance of striped data also increases.

Stripe size

The size of the stripes written to each disk. This may also be referred to as block size, chunk size, stripe length, or granularity.

You should use a large stripe size of at least 1 MB.

If you are using RAID devices, you should use RAID 10 or RAID 01 because it offers the best performance of all RAID systems and good fault tolerance.

Selecting the Analytics server host

When selecting a host to be the Platform Analytics server, you need to ensure that the host is running a supported operating system. See *IBM Platform Analytics Release Notes* for the latest list of supported operating systems for the Analytics server host.

Tip: If you select a host that also meets the Tableau Server system requirements, you can also select the Platform Analytics server host to be the Platform Analytics reporting server host.

For optimal performance, the Analytics server host should be a dedicated multi-core host with sufficient memory and input/output performance. The network bandwidth between the Analytics server, the database hosts, and the Analytics nodes is a key performance factor in the Analytics server host.

If you are not using asynchronous data loading mode, the hardware configuration shown in Table 5 should be sufficient.

Table 5. Recommended hardware specifications for the Analytics server host, not using asynchronous data loading

RAM	CPU	Local hard disk	Network
4 GB	4 × 2.4 GHz	7200 RPM SATA/SCSI/SAS 50 GB	Gigabit Ethernet

If you are using asynchronous data loading mode, memory is a key performance factor for the Analytics server host. If the Analytics server is running on a Microsoft Windows host, you should use the 64-bit version of Windows because Java™ cannot use more than 1638MB of memory on 32-bit platforms.

You should only use the asynchronous data loading mode for sending data from the Analytics node to the database over a slow or unstable network.

Table 6 describes the optimal hardware configuration of the Analytics server if you are using asynchronous data loading, depending on the size of your cluster.

Table 6. Recommended hardware specifications for the Analytics server host, using asynchronous data loading

Cluster size	RAM	CPU	Local hard disk	Network
Medium (100 - 1000 hosts)	4 GB	4 × 2.4GHz	7200 RPM SATA/SCSI/SAS 50 GB	Gigabit Ethernet
Large (more than 1000 hosts)	8 GB	4 × 2.4GHz	7200 RPM SATA/SCSI/SAS 50 GB	Gigabit Ethernet

Note:

Using an NFS disk mount instead of a local hard disk is not recommended.

Selecting the Analytics reporting server host

When selecting a host to be the Analytics reporting server host, you must ensure that it meets the detailed system requirements for Tableau Server shown in Table 7.

Table 7. Hardware configuration requirements for Tableau Server

Operating system	RAM	CPU	Services	User accounts
Windows Server 2003 (SP2 or higher) Windows Server 2008 Windows Server 2008 R2	4 GB and above	Dual-core	Do not run Internet Information Services (IIS) to avoid conflicts with the web server port 80.	Access to an administrator account to install software and services. Access to a user account that the service can use (optional).

Tip: If the Analytics server host also meets the Tableau Server system requirements, you can select the Analytics server host to also be the Platform Analytics reporting server host.

Refer to the Tableau Server documentation or to *IBM Platform Analytics Release Notes* for the latest list of system requirements and supported operating systems for the Tableau Server.

The network bandwidth between the Platform Analytics reporting host and the database cluster may be an important performance bottleneck. Therefore, the Platform Analytics reporting host should have a Gigabit Ethernet connection with database hosts.

Table 8 on page 10 describes the optimal hardware configuration of the Platform Analytics reporting server, depending on the size of your cluster.

Table 8. Recommended hardware specifications for the Analytics reporting server host

Cluster size	RAM	CPU	Local hard disk	Network
Medium (100 - 1000 hosts)	4 GB	4 × 2.4GHz	7200 RPM SATA/SCSI/SAS 50 GB	Gigabit Ethernet
Large (more than 1000 hosts)	8 GB	4 × 2.4GHz	7200 RPM SATA/SCSI/SAS 50 GB	Gigabit Ethernet

Selecting the Analytics node hosts

When selecting a host in the cluster to be a Platform Analytics node, you need to ensure that the host is running a supported operating system, and that it meets the minimum hardware requirements. See *IBM Platform Analytics Release Notes* for the latest system requirements for the Analytics node host.

The Analytics node must run on a management host in each Platform Symphony cluster from which you want to collect data.

For optimal performance of your Analytics node, the host should be running on a high-bandwidth network. Since network bandwidth is an important performance bottleneck for the Analytics nodes, the Analytics node host should have a Gigabit Ethernet connection with the database host.

Table 9 describes the optimal hardware configuration of the Analytics node based on the size of the clusters in which the node resides.

Table 9. Recommended hardware specifications for the Analytics node hosts

Cluster size	RAM	CPU	Local hard disk	Network
Medium (100 - 1000 hosts)	4 GB	2 × 2.4GHz	7200 RPM SATA/SCSI/SAS 50 GB	Gigabit Ethernet
Large (more than 1000 hosts)	8 GB	4 × 2.4GHz	7200 RPM SATA/SCSI/SAS 50 GB	Gigabit Ethernet

Preparing the LSF clusters (LSF 7.x, 8.x and up only)

Use the following procedure to prepare the LSF clusters.

Before you begin

Your clusters must be running one of the following versions of Platform LSF:

- LSF 7.x
- LSF 8.0
- LSF 8.0.1
- LSF 8.3
- LSF 9.1

You can skip this step for any clusters that are not running any of these versions of Platform LSF.

About this task

By default, LSF 7, 8, and 9 do not enable streaming of `lsbatch` system events. If you want to collect data from LSF stream files, you must enable streaming of `lsbatch` system events.

Procedure

1. Log into a host in the LSF cluster.
2. Edit the `lsb.params` file.
   - UNIX: `$LSF_ENVDIR/lsbatch/`*`cluster_name`*`/configdir/lsb.params`
3. In the `lsb.params` file, edit the **Parameters** section to enable the exporting of LSF job event data to the `lsb.stream` file.

   Add the following lines to the **Parameters** section:

   ```
   # Enable streaming of lsbatch system events
   ENABLE_EVENT_STREAM=y
   # Determines the location of the lsb.stream file. This parameter is optional.
   # The default location is: $LSB_SHAREDIR/{clustername}/logdir/stream.
   # EVENT_STREAM_FILE=/tmp/lsb.mystream
   # Determines the maximum size of the lsb.stream file. This parameter is optional.
   # The default size is 100MB.
   # MAX_EVENT_STREAM_SIZE=1024
   ```

   Note: Also edit the following parameters if you have LSF 8.x

   ```
   # The ALLOW_EVENT_TYPE parameter prevents unwanted events in the stream files. Only
   defined events in this parameter are logged into the stream series of files, lsb.stream,
   lsb.status and lsb.pendingreasons. The value of this parameter is a collection of LSF
   event types. The event type string is the same as those in the lsb.events file. As a
   minimum, you have to define the following events in the ALLOW_EVENT_TYPE to provide
   sufficient information for Analytics to generate reports.
   ALLOW_EVENT_TYPE=JOB_NEW JOB_FINISH2 JOB_STARTLIMIT JOB_STATUS2 JOB_PENDING_REASONS
   # This parameter defines the interval for mbd to fork a child to log the pending and
   running jobs. The suggested value is 15 minutes, and if it is not defined, the default
   behavior is that the parameter is turned off.
   RUNTIME_LOG_INTERVAL=10
   ```

4. If **ENABLE_EVENT_STREAM** is turned on and **ALLOW_EVENT_TYPE** includes `JOB_PENDING_REASONS`, configure the following parameters to control the streaming of job pending events to the `lsb.pendingreasons` file:

   ```
   #GROUP_PEND_JOBS_BY = field_name & field_name &
   #Available fields for the above parameter are (case insensitive):
   #QUEUE,USERNAME,RES_REQ,PROJECT_NAME,USER_GROUPS,APPLICATION
   #JOB_GROUP,NUM_PROCESSORS,LICENSE_PROJECT,HOST_TYPE

   #PENDING_TIME_RANKING = short[1,20] medium[21,400] long[401,]
   #CONDENSE_PENDING_REASONS = y
   #INCLUDE_DETAIL_REASONS = y
   ```

   where:

   GROUP_PEND_JOBS_BY
   Defines the key fields, other than the pending reasons, used to group jobs. You can use this parameter to limit the number of pending event records logged. The more keys you include, the more records will be logged in `lsb.pendingreasons`. If this parameter is not configured, the default is to group pending jobs by pending reasons and pending time ranking. If there is an error in the configuration of this parameter, then grouping will be done by `USERNAME & QUEUE`.

   PENDING_TIME_RANKING
   Defines and classifies pending jobs by the time they have been pending, in minutes. You can use this ranking to identify long-pending jobs, while not losing information about short-pending jobs. A job meets the criteria for a

given rank if its pending time is greater than the lower boundary and less than or equal to the upper boundary for that rank. The boundaries for a given rank must not overlap the boundaries of another rank, but gaps between ranks are allowed.

CONDENSE_PENDING_REASONS
Controls whether or not host-based pending reasons are condensed into one generic pending reason. When this is set to `y`, **mbatchd** will not get detailed pending reasons from the scheduler. You can use the **INCLUDE_DETAIL_REASONS** parameter to override this.

INCLUDE_DETAIL_REASONS
Controls whether or not the scheduler sends detailed pending reasons to **mbatchd** for logging when **CONDENSE_PENDING_REASON** is turned on. The default value is `y`.

5. Reconfigure **mbatchd** to apply these changes:
   ```
   badmin mbdrestart
   ```
6. To verify that these changes are in effect, verify that the `lsb.stream` file exists.
   By default, `lsb.stream` is located at the following directories:
   - UNIX: `$LSB_SHAREDIR/cluster_name/logdir/stream`

   If you defined the **EVENT_STREAM_FILE** parameter in `lsb.params`, check the specified file path for the `lsb.stream` file.

Obtaining the installation files

Use the following procedure to obtain the installation files for Platform Analytics, the Analytics database, and the Tableau Server.

- Obtain the necessary files for installing Platform Analytics.
  You need the following files to install Platform Analytics:
  - IBM Platform Analytics server installation package
  - IBM Platform Analytics node installation package
  - IBM Platform Analytics documentation package
- Obtain the necessary files for installing the Analytics database.
  You need the following files to install the Analytics database:
  - Analytics database installation package
  - IBM Platform Analytics database schema package
- Obtain the necessary files for installing Tableau Server.
  You need the following files to install the Tableau Server:
  - Tableau Server installation package

Pre-installation checklist

Use this pre-installation checklist prior to installing Platform Analytics.

System sizing depends on a number of factors, the most significant include:

- Number of jobs per day across all clusters
- Number of hosts to be monitored
- Number of clusters to be monitored
- Number of concurrent logins to be supported
- Amount of job data to be retained
- Job graph archiving choice

Consider these factors, along with the specifications listed in "Selecting the hosts" on page 7, to help plan your installation.

Keep the following tips in mind before you install Platform Analytics:

- Database sizing depends on the number of jobs per day and job detail data retention.
- Ensure that the Analytics database data directory path is the same on each node.
- Make sure that you follow the security setting that is required for the Analytics database installation, as described in the installation guide for the Vertica Analytic Database.
- Make sure that the Analytics binary that you have downloaded matches the operating system and the system architecture on which you intend to install it.
- Verify that you can ping the database nodes from the Analytics node, Tableau Server, and Analytics server.
- Verify that the Analytics node can ping the Analytics server and vice versa.
- Make sure that the Analytics server and Tableau Server meets hardware configuration requirements. If you are planning to install both on the same system, then make sure your hardware supports and withstands both configurations.
- Verify that you have enough memory and disk space for each node.

Chapter 3. Installing Platform Analytics

After preparing your hardware, you can install the Analytics database, the Analytics server, and the Analytics node. Then, review the items on the post-installation checklist.

Installing the Analytics database

The Analytics database, which consists of one or more database hosts, contains the cluster operations data for reporting and analysis. You must install and configure the Analytics database before you install Platform Analytics.

You can install the Analytics database on a single host or on a cluster of multiple hosts. Installing the database on multiple hosts improves the performance and reliability of your database.

Before you begin

The intended database host (or hosts in the cluster) must meet the Platform Analytics and Vertica system requirements.

Manually check the following items for the Analytics database cluster hosts:

- Verify that the operating system type and version of the cluster hosts meet the Vertica requirements. All hosts in the database cluster should be running the same operating system type and version.
- Ensure that the hosts have at least 2 GB of swap space.
- Verify that the configured data and catalog directories have been created on all cluster hosts.
- Verify the Python version; only version 2.x, higher than 2.3.4, is supported.
- Verify that the rsync version is 3.0.5 or higher.
- Verify the ssh access configuration.
- Verify that NTP is running on startup.
- Configure host name resolution.
- Define the loopback address.
- Ensure that port 5433 is not in use and that `/dev/pts` is mounted.

Manually check the following items as described in the Vertica documentation, *Installation and Configuration Guide*, before installing the database:

- Make a note of disk storage locations.
- Make note of IP addresses of the hosts in the database cluster.
  - If you are using a host with multiple network interfaces, use the IP address assigned to the NIC that is connected to the other cluster hosts, not the NIC that is used for client connections.
  - Dynamically assigned IP addresses are not supported.
  - Subnet Masks must be the same for all hosts.
- Set the host locale (Language).
- Set the default time zone.
- Remove nonessential applications.

- Disable firewall.
- Optionally, run Spread on a separate control network.
- Back up any existing database.
- Provide root SSH access to all the cluster hosts (with the same password).

For details about installing the database, see the Vertica documentation, *Installation and Configuration Guide*. Follow the steps described in the "Before You Install" and "Installing Vertica" chapters.

Procedure

1. Log in to the intended database host (or hosts in the cluster) as **root**.
2. Extract the installation package to a shared directory that all the cluster hosts can access.
3. Edit the `install.config` file and configure all required parameters.

   You must configure all the required parameters in the `install.config` file. To configure any required or optional parameter, uncomment the line and specify the appropriate value.
4. Run the database installation script:

   ```
   ./AnalyticsVerticaInstaller.sh -f install.config
   ```

   The installer performs the following steps:

   a. Initializes configurations for the installation
   b. Installs/repairs SSH keys for the **root** account on the cluster hosts
   c. Checks prerequisites on the cluster hosts

      Note: If there are any errors in the previous three steps, then an `install.log` file is generated in the installation folder. Fix the errors and rerun the installation script.
   d. Installs and configures the Analytics database
   e. Creates the database connection for Platform Analytics
   f. Deploys the Platform Analytics database schema

      Note: If there are any errors in the previous three steps, then an `install.log` file is generated in the installation folder. Fix the errors, then perform the following steps:

      1) Uninstall the database rpm package by running this command:

         ```
         rpm -e vertica-<vertica version>
         ```

      2) Delete the `/opt/vertica` directory, and the files and directories under your `DB_DATA_PATH` and `DB_CATALOG_PATH` paths.
      3) Rerun the database installation script.

What to do next

If you need to run the **dbconfig** script to update any of the parameters in the database `install.config` file, you must first source the Analytics environment before running **dbconfig**:

- For **csh** or **tcsh**:

  ```
  source ANALYTICS_TOP/conf/cshrc.perf
  ```

- For **sh**, **ksh**, or **bash**:

  ```
  . ANALYTICS_TOP/conf/profile.perf
  ```

Installing the Analytics server

Perform this task to install the Platform Analytics server on a host.

Before you begin

Check to ensure the following items:

- The Analytics server meets the detailed system and software requirements.
- The Analytics server's operating system can handle daylight savings time correctly.
- Removed any previous installations of Platform Analytics from your host.
- The required ports for Platform Analytics and database hosts are free. If a firewall exists, all these ports must be open in the firewall. All the required ports are TCP.
- The Analytics database host is properly configured and running.
- The Analytics server has access to the Analytics database JDBC driver.

Procedure

1. Install and configure Tableau Server. Make sure that the `Tableau Server (tabsvc)` service has started.

   Note: You will find a separate installer for Tableau Server in the installation package. Refer to the Tableau Server documentation if you need more installation details.
2. Log in to the intended Platform Analytics server as administrator.
3. Extract the Platform Analytics server package and edit the `install.config` file.

   Important: Configure all the required parameters in the `install.config` file. To configure any required or optional parameter, uncomment the line and specify the appropriate value.
4. Navigate to the top-level installation directory and run the Platform Analytics server installation package:

   Note: Make sure you have administrator privileges before running this batch file.

   ```
   AnalyticsServerInstaller.bat -f install.config
   ```

   Note:

   - You can receive notifications of events via email. This means you will be made aware of potential problems in your cluster without having to constantly monitor the Platform Analytics Console.
   - After enabling email notifications, specify the SMTP mail server and the intended recipient of the email notifications. The intended sender is the account that is used to send the email notifications.
5. Verify that the **pats** service has started:
   - Windows: *ANALYTICS_TOP*`\bin\perfadmin list`

Installing the Analytics node

Install the Platform Analytics node package on a host. You must have one Analytics node host in each cluster that Platform Analytics will handle.

Before you begin

Check to ensure the following items:

- The Analytics node host meets the detailed system and software requirements.
- The Analytics node host operating system can handle daylight savings time correctly.
- If you want the Analytics node to collect LSF cluster data, check the following:
  - You have access to the primary LSF administrator account.
  - The Analytics node host must be an LSF server in the cluster with access to the LSF event file (`lsb.stream`).
- The required port for the Platform Analytics hosts is free. If a firewall exists, this TCP port must be open in the firewall.
- The Analytics database host is properly configured and running.
- The Analytics node host has access to the Analytics database JDBC driver.

Procedure

1. Log in to the intended Analytics node host. If you want to update the `rc.boot` scripts, log in as **`root`**.
2. Extract the Analytics node package and edit the `install.config` file.

   Note: Configure all the required parameters in the `install.config` file. To configure any required or optional parameter, uncomment the line and specify the appropriate value.
3. Run the Analytics node installation script:

   ```
   ./AnalyticsNodeInstaller.sh -f install.config
   ```

   Or, for silent installation:

   ```
   ./AnalyticsNodeInstaller.sh -f install.config -silent
   ```
4. Source the PERF environment.
   - For **`csh`** or **`tcsh`**:

     ```
     source ANALYTICS_TOP/conf/cshrc.perf
     ```
   - For **`sh`**, **`ksh`**, or **`bash`**:

     ```
     . ANALYTICS_TOP/conf/profile.perf
     ```

   In this documentation, *`ANALYTICS_TOP`* represents the full path to the top-level installation directory for the Analytics node on your host.

   Note: After sourcing the PERF environment on the Red Hat Enterprise Linux 6.*x* x_64 platform, you might notice an error, such as:

   ```
   -bash: /lib/libc.so.6: No such file or directory
   -bash: [: : integer expression expected
   Cannot get binary type
   ```

   To avoid this error, after sourcing the PERF environment, you must also install the related files by running the following command:

   ```
   yum install ld-linux.so.2
   ```
5. If you install Platform Application Center in the same cluster as Platform Analytics, disable the auto delete function of LSF events loader (**`lsfeventsloader`**) from Platform Application Center to prevent the `lsb.stream.UTC` files from being deleted automatically.
   a. Edit the `$PERF_CONFDIR/conf/dataloader/lsbevents.properties` file of Platform Application Center.
   b. Add the following line:

```
AUTO_DELETE_STREAM_FILE=N
```

6. Manually stop and start the **plc** service:

```
perfadmin stop plc
perfadmin start plc
```

7. Verify that the data loaders are running:

```
plcclient.sh -s
```

8. Verify that there are no errors in any of the data loader log files.

   The data loader log files are located in the `dataloader` subdirectory of the PERF log directory, `$PERF_LOGDIR/dataloader`.

Post-installation checklist

After installation, verify that Platform Analytics is installed successfully and working properly.

Use the following checklists to verify each part of your Platform Analytics installation.

Analytics database

- Ensure that the nodes are up and running.
- Verify that you can create a database user for the creation of Platform Analytics schema.
- Make sure that the **K_SAFE** parameter in the create schema file is set to the correct value. Use the following command to check the value of the **K_SAFE** parameter:

  ```
  SELECT current_fault_tolerance FROM system;
  ```

- Ensure that the network bandwidth utilization is set properly.

Analytics node

- If you are collecting LSF cluster data, source the LSF environment.
- Start the Analytics node after you start the Platform Analytics server.
- Wait for 10 minutes after you install the Analytics node and then check the `lsfeventsloader`, `lsfbhosts`, and `hostmetrics` log files to make sure data has been written and there are no issues for each node.
- Verify that the database is receiving data. If you have configured remote data loading, then check the log file to make sure data are written correctly.

Note: If you later upgrade Platform Symphony after having installed Platform Analytics, you must either change the value of **SOAM_VERSION** in the `perf.conf` file to match the new Symphony version, or reinstall the Analytics node after upgrading Symphony.

Analytics server

- Verify that you can connect to the Analytics database.
- When the server runs two data transformers, check the log files. Also log on to the Analytics database to verify that the **aggregate_info** has the proper date and proper tables updated.
- Verify that the report table contains data:

  ```
  select count(*) from rpt_<data transformer name>_raw;
  ```

  where *data transformer name* can be **jobmart**, **workload_statistics**, **cluster_capacity**, **hardware**, or **flexlm_licusage**.

- After you install the Analytics server and Analytics node, wait two hours and then check the **Cluster** tab to see if all hosts are added.
- Make sure you enable event notification during the Analytics server installation and start the server before you start the Analytics node.

  Check the `event_manager_conf` table to make sure that your event manager is set to active:

  ```
  event_manager_conf.EVENT_MGR_ACTIVE ='Y'
  ```
- Verify that email notification is set up correctly.

Tableau Server

- Verify that you can activate the server. If offline activation is required, then proceed with the on-screen instructions.
- If you are not using port 80 during the install configuration, write down the port number.
- Make sure you remember the Administrator user name and password.
- When you log in as Administrator, make sure you select **Embedded Credentials** on the Maintenance page.
- Verify that the tableau service is up and running on the Maintenance page.
- Verify that the window service for the Tableau Server is set to `Automatic` and has started.

Platform Analytics reporting server

- Verify that you have extracted the reports package on the Tableau Server.
- Verify that the deployment is successful and the workbooks are under the projects folder.
- Make sure that the `parb` service starts from Windows services console.

Platform Application Center plug-in (optional)

- Verify that the report daemon is up and running.
- Verify that the `report.conf` file is configured correctly and that you can ping the host that has Platform Application Center installed.
- Verify that you can successfully schedule a report and send a report.

Chapter 4. Configuring and optimizing Platform Analytics

After you install Platform Analytics, you can verify and configure the Analytics server, configure the Analytics node, and optimize the Analytics server, node, and database.

Verifying and configuring the Analytics server

Use the IBM Platform Analytics Console to verify and configure the Analytics server.

Procedure

1. Log in to the Analytics server host.
2. Launch the Platform Analytics Console.
   - Windows: **Start** > **All Programs** > **>IBM Corporation** > **Analytics Server** > **Analytics Console**
3. Click **Data Collection Node** in the navigation tree and verify that the node is running correctly.

   To view the data loader properties, right-click each loader controller instance and select **Loader Properties**.
4. Click **Scheduled Tasks** in the navigation tree and modify the times that the scheduled tasks are running, if necessary.
5. Click **Events in the navigation tree** and verify that there are no `ERROR` or `FATAL` events.
6. Verify the email notification settings.

   While in **Events**, click **Action** > **Notification** to open the **Event Notification** dialog.

Enabling and configuring the remoting server (asynchronous data loading mode only)

If you intend to use asynchronous data loading mode, perform this task to enable and configure the remoting server.

About this task

You should only use the asynchronous data loading mode for sending data from the Analytics node to the database over a slow or unstable network.

Normally, data loaders perform synchronous data loading, whereby they load data directly into the Analytics database. In rare cases where the network connection between the Analytics node and the database host is poor, the data loaders will perform asynchronous data loading. In such cases, the data loaders send data to the Analytics server, and the server then loads the data into the Analytics database.

Procedure

1. Navigate to the *ANALYTICS_TOP*`/conf/` directory.
2. Rename the `remotingserver.TMPL` file to `remotingserver.xml`.
3. Edit the `remotingserver.xml` file and change the IP address to the IP address of the Analytics server host.

4. Edit the `log4j.properties` file and uncomment the `RemotingServer` line.
5. Rename the `wsm/pars.TMPL` file to `wsm/pars.conf`.
6. Start the remoting server:

   ```
   perfadmin start pars
   ```
7. If you are on a Windows host and want the remoting server to start automatically when Windows starts, change the startup type of the **pars** Window service from **Manual** to **Automatic**.

Modifying the LSF version if you have LSF 8 but have not applied the LSF 8 add-on data solution

Follow these steps only if you have LSF 8 installed and you have not applied the LSF 8 add-on data solution (`Oct Qpk`).

Procedure

1. Log in to the Analytics server host.
2. Edit the `$ANALYTICS_TOP/conf/pi.conf` file and navigate to the following line:

   ```
   LSF_VERSION = version
   ```
3. Change *version* from `8` to `7`, only if it is currently set to `8`.

   By default, *version* is set to `7`.

   Note: If you do not change the version, the data you see may not be accurate.

Enabling support for collecting LSF XL data

Procedure

Perform this task if you want to collect data from LSF XL.

1. Log in to the Analytics server host.
2. Edit the *ANALYTICS_TOP*`/conf/pi.conf` file, locate the **SUPPORT_LSF_XL** parameter, and change the value from `N` to `Y`:

   ```
   SUPPORT_LSF_XL=Y
   ```

Configuring the Analytics node

You can perform the following tasks to configure the Analytics node to work more efficiently with your cluster.

Stopping redundant PERF services (LSF 7 clusters only)

Platform Analytics runs its own PERF services independent of the cluster. Since LSF 7 clusters also have PERF services running, you should stop the redundant PERF services to avoid unnecessary redundancy in your cluster.

Before you begin

You can only stop the PERF services if you installed the Analytics node in a cluster running LSF 7.

Procedure

1. Log in to any host in the LSF cluster as the cluster administrator.
2. If you logged into a UNIX host, source the LSF environment.
   - For **csh** or **tcsh**:

     `source LSF_TOP/conf/cshrc.lsf`
   - For **sh**, **ksh**, or **bash**:

     `. LSF_TOP/conf/profile.lsf`
3. Navigate to the PERF binary directory.
   - UNIX: `cd $PERF_TOP/version_number/bin`
4. Stop the loader controller (**plc**) and data purger (**purger**) services:

   ```
   perfadmin stop plc
   perfadmin stop purger
   ```
5. Rename the `plc.xml` loader controller configuration file to another file extension to remove it from the list of LSF services. For example, rename `plc.xml` to `plc.xml.bak`.
6. Rename the `purger.xml` data purger configuration file to another file extension to remove it from the list of LSF services. For example, rename `purger.xml` to `purger.xml.bak`.
7. Restart EGO on the LSF master host to activate these changes:

   `egosh ego restart lsf_master_host_name`

Modifying the LSF version if you have LSF 8 but have not applied the LSF 8 add-on data solution

Follow these steps only if you have LSF 8 installed and if you have not applied the LSF 8 add-on data solution (`Oct Qpk`).

Procedure

1. Log in to the Analytics node host.
2. Edit the `$PERF_TOP/conf/perf.conf` file and navigate to the following line:

   LSF_VERSION = 8.0
3. Change `8.0` to `7.0`.

   If you do not change the version, the data you see may not be accurate.

   Note: Defining **LSF_VERSION** in an Analytics node as `7.0` changes the following loader behaviors:
   - **lsfeventsloader** loads **JOB_FINISH** events from `lsb.stream` files instead of **JOB_FINISH2** events.
   - Data loading of **bjobs** data is sampled by **lsfbjobloader** instead of **lsfjobstatusloader**.

- Data loading of jobs pending reason is sampled by **lsfpendingreasonloader** instead of **lsbpendingreasonloader**.

4. Navigate to the LSF directory, `$LSF_LIBDIR/../../../`, and create a linked directory, `7.0`:

   ```
   ln -s 8.0 7.0
   ```

   Note: This is for Platform Analytics data loaders to dynamically link to LSF libraries after changing **LSF_VERSION** defined in the Analytics node.

Configuring the LSF cluster for the host core utilization data loader

Configure **elim** in the LSF cluster to allow the **hostcoreutilloader** data loader to collect data from the cluster.

Procedure

1. Log in to the Analytics node host.
2. If you logged in to a UNIX host, source the LSF environment.
   - For **csh** or **tcsh**:

     ```
     source LSF_TOP/conf/cshrc.lsf
     ```

   - For **sh**, **ksh**, or **bash**:

     ```
     . LSF_TOP/conf/profile.lsf
     ```

3. Copy the `elim.coreutil` execution files from the Platform Analytics directory to the relevant LSF directory.
   - UNIX:

     ```
     cp ANALYTICS_TOP/elim/os_type/elim.coreutil $LSF_SERVERDIR
     ```

     where *os_type* is either `aix`, `hpux`, `linux`, or `solaris`, depending on the specific UNIX operating system.

4. Edit the *LSF_TOP*`/conf/lsf.shared` file and add **CORE_UTIL** as a resource.

   Add the following line to the **Resource** section:

   ```
   CORE_UTIL     String   300     ()          (Core Utilization)
   ```

   For example:

   ```
   Begin Resource
   RESOURCENAME  TYPE    INTERVAL INCREASING  DESCRIPTION
   CORE_UTIL     String   300     ()          (Core Utilization)
   End Resource
   ```

5. Edit the *LSF_TOP*`/conf/lsf.cluster.`*cluster_name* file and add **CORE_UTIL** as a resource map.

   Add the following line to the **Resource** section:

   ```
   CORE_UTIL     [default]
   ```

   For example:

   ```
   Begin ResourceMap
   RESOURCENAME  LOCATION
   CORE_UTIL     [default]
   End ResourceMap
   ```

6. Reconfigure the LSF cluster to apply your changes:

   ```
   badmin reconfig
   lsadmin reconfig
   ```

Enabling FLEXnet data loaders

Perform this task to enable the FLEXnet data loaders in your cluster.

Before you begin

Check the following items:

- IBM Platform Analytics is configured to collect FLEXnet license data.
- The FLEXnet usage data loader and the FLEXnet license server are using the same time zone.

Procedure

1. Log in to the Analytics node host.
2. Modify the FLEXnet usage data loader configuration file to set up a server list pointing to the FLEXnet license servers:
   a. Edit the FLEXnet usage data loader configuration file.
      - UNIX: *ANALYTICS_TOP*/conf/dataloader/flexlicusage.properties
   b. Change the **ServerList** (or **FileName**) parameter to point to the FLEXnet license servers.
3. Modify the FLEXnet events data loader configuration file to point to the license log file:
   a. Edit the FLEXnet events data loader configuration file.
      - UNIX: *ANALYTICS_TOP*/conf/dataloader/flexlicevents.properties
   b. Change the **LicenseLogFile** parameter to point to the FLEXnet license log file. For example:
      - UNIX: LicenseLogFile=/*file_path*/lmgrd.log
4. Start the Analytics node host by restarting the loader controller (plc) service:

```
perfadmin stop plc
perfadmin start plc
```

Enabling FLEXnet Manager data loaders

Perform this task to enable the FLEXnet Manager data loaders in your cluster.

Before you begin

You can only enable FLEXnet Manager data loaders if you configured Platform Analytics to collect FLEXnet license data.

Procedure

1. Log in to the Analytics node host.
2. If you connected to the UNIX host via **telnet** and are running **xserver** on a local host, set your display environment.

   Test your display by running **xclock** or another X-Windows application.

   If the application displays, your display environment is already set correctly; otherwise, you need to set your display environment.

   - For **csh** or **tcsh**:

     ```
     setenv DISPLAY hostname:0.0
     ```

   - For **sh**, **ksh**, or **bash**:

     ```
     DISPLAY=hostname:0.0
     export DISPLAY
     ```

   where *hostname* is your local host.
3. Add a data source for the FLEXnet Manager Reporting database to the Analytics node.

   The FLEXnet Manager Reporting data source is named FNMReportDB.

- In UNIX, run:

  ANALYTICS_TOP/1.2/bin/dbconfig.sh add FNMReportDB

4. Modify the FLEXnet Manager data loader configuration file to point to the FLEXnet Manager servers:
   a. Edit the FLEXnet Manager data loader configuration file.
      - UNIX: *ANALYTICS_TOP*/conf/dataloader/fnmloader.properties
   b. Change the **DataSource** parameter to specify the name of the data source for the FLEXnet Manager Reporting database.

      Each FLEXnet Manager server has its own data source, and each data loader can only access one data source; therefore, for each FLEXnet Manager server, you need to specify a separate data source for each data loader.

      For example, for the FLEXnet Manager Reporting database:

      ```
      DataSource=FNMReportDB
      ```
5. Enable the FLEXnet Manager data loader in your cluster:
   a. Edit the loader controller configuration file for the FLEXnet Manager data loader.
      - UNIX: *ANALYTICS_TOP*/conf/plc/plc_license.xml
   b. Enable data gathering for the FLEXnet Manager data loader (fnmloader) by modifying the **Enable** attribute of the <DataLoader Name="fnmloader" .../> element.

      The element should now resemble the following:

      ```
      <DataLoader Name="fnmloader" ... Enable="true" ... />
      ```
6. Start the Analytics node host by restarting the loader controller (plc) service:

   ```
   perfadmin stop plc
   perfadmin start plc
   ```

Installing and configuring the FLEXnet Manager scripts

Perform this task to install and configure the Platform FLEXnet Manager (FNM) scripts to work with the FLEXnet Manager data loaders and your LSF cluster.

Before you begin

You can only configure the FLEXnet Manager scripts if you configured Platform Analytics to collect FLEXnet license data.

Check that the following are installed and functional:

- FLEXnet Manager, version 11 or later
- Jasper reports

Procedure

1. Enable the FLEXnet Manager data loader:
   a. Edit the loader controller configuration file for license data loaders.

      Edit *PERF_TOP*/conf/plc/plc_license.xml.
   b. Enable the FLEXnet Manager data loader.

      ```
      <DataLoader Name="fnmloader" ... Enable="true" ... />
      ```
   c. Restart the **plc** service for your changes to take effect.
2. Deploy the **esub** script.

This script provides the cluster name and job submission time to the **LM_PROJECT** environment variable, which passes the data to FLEXnet Manager to be collected by the Platform Analytics data loaders.

a. If you have an existing script in place, add the following information to it:

```
modenv(LM_PROJECT => substr($PROJECT_NAME, 0, 5).",$LSF_CLUSTER_NAME,$SUBMIT");
```

Note the following:

- **$LSF_CLUSTER_NAME** is the name of the LSF cluster where the job is running
- **$SUBMIT** is the job submission time as the number of non-leap seconds from 00:00:00 UTC, January 1, 1970.

b. If you do not have an existing script, copy and deploy the `esub.fnm` script, which is located in the PERF samples directory:

- UNIX: `$PERF_TOP/samples`

3. Deploy the job starter script.

This script provides the job ID and job array index to the **LM_PROJECT** environment variable, which passes the data to FLEXnet Manager to be collected by the Platform Analytics data loaders.

a. Edit the `lsb.queues` file and add the **JOB_STARTER** parameter with the path to the **jstart** script to each queue that you want to control a submitted.

```
JOB_STARTER = /path_to_jstart/jstart
```

b. In the command console, reconfigure the master host to activate this change.

```
badmin reconfig
```

c. In the command console, display detailed queue information to verify the configuration.

```
bqueues -l queue_name
```

For example, if you added the script to the Normal queue:

```
bqueues -l normal
```

d. If you have an existing script in place, add the following information to it:

```
LM_PROJECT=$LM_PROJECT,$LSB_JOBID,$LSF_JOBINDEX
export LM_PROJECT
```

e. If you do not have an existing script, copy and deploy the **jstarter.fnm** script, which is located in the PERF samples directory:

- UNIX: `$PERF_TOP/samples`

4. Deploy the FNM data loader to a data collection host.

Enabling and configuring the remoting client (asynchronous data loading mode only)

If you intend to use the asynchronous data loading mode, perform this task to enable and configure the remoting client.

About this task

You should only use the asynchronous data loading mode for sending data from the Analytics node to the database over a slow or unstable network.

Normally, data loaders perform synchronous data loading, whereby they load data directly into the Analytics database. In rare cases where the network connection between the Analytics node and the database host is poor, the data loaders will perform asynchronous data loading. In such cases, the data loaders send data to the Analytics server, and the server then loads the data into the Analytics database.

Procedure

1. Log in to the Analytics node host.
2. Navigate to the *ANALYTICS_TOP*/conf directory.
3. Rename the remotingclient.TMP file to remotingclient.xml.
4. Edit the remotingclient.xml file and change the **serverBindAddress** and **serverBindPort** parameters according to the remoting server configuration.
5. Edit the perf.conf file and add the following parameters:

   ```
   ASYNC_LOADING_ENABLED=Y
   ASYNC_LOADING_MODE=Remoting
   ```

6. Restart the loader controller service (**plc**).

   ```
   perfadmin stop plc
   perfadmin start plc
   ```

Enabling support for collecting LSF XL data

Procedure

Perform this task if you want to collect data from LSF XL.

1. Log in to the Analytics node host.
2. Edit the *ANALYTICS_TOP*/conf/perf.conf file and add the following line to enable LSF XL support:

   ```
   SUPPORT_LSF_XL=Y
   ```

Optimizing the Analytics server

You can perform the following tasks to modify the Analytics server to enhance performance.

About this task

The following tasks are optional.

Procedure

1. "Changing the data retention period"
2. "Splitting data transformer tasks to disperse workload" on page 30

Changing the data retention period

You can perform this task to change the data retention period.

About this task

A long data retention period can have a significant impact on performance and the data volume. You can tailor the data retention period according to your business requirements to maximize the performance of your Analytics server.

The data purger consists of multiple scheduled tasks (**PartitionMaintenanceGroup***), which are enabled by default.

Procedure

1. Launch the **vsql** command line.
   a. Navigate to the bin subdirectory of the database installation directory. By default, this is /opt/vertica/bin.

b. Run **vsql** to connect to the database:

   ```
   ./vsql -d database_name -p port -U username -w password
   ```

   where:

   - *database_name* is the name of the database
   - *port* is the TCP port number or the local socket file extension in which the server is listening for connections. The default is port number 5433.
   - *username* is the name of the user with which to connect to the database, instead of the default user (the database administrator).
   - *password* is the password for the database user.

   Alternately, you can run **vsql** with no options to accept the defaults and specify the administrator password at the prompt.

2. Examine the current data retention periods of the database tables for Platform Analytics.

   - To examine the retention periods for all Analytics database tables, run the following from the **vsql** command line:

     ```
     SELECT TABLE_NAME, DATA_DAYS_RANGE
     FROM SYS_TABLES_TO_PARTITION;
     ```

   - To examine the retention periods for a specific database table, run the following from the **vsql** command line:

     ```
     SELECT TABLE_NAME, DATA_DAYS_RANGE
     FROM SYS_TABLES_TO_PARTITION
     WHERE TABLE_NAME='table_name';
     ```

     where *table_name* is the name of the table you want to examine.

   The output displays the name of the table and the corresponding data retention period in days.

   For example, to view the data retention period for the RESOURCE_METRICS_BUILTIN table, run the following from the **vsql** command line:

   ```
   SELECT TABLE_NAME, DATA_DAYS_RANGE
   FROM SYS_TABLES_TO_PARTITION
   WHERE TABLE_NAME='RESOURCE_METRICS_BUILTIN';
   ```

3. Change the data retention period for the appropriate database tables.

   a. For each database table to change, run the following from the **vsql** command line:

      ```
      UPDATE SYS_TABLES_TO_PARTITION
      SET DATA_DAYS_RANGE='retention_period'
      WHERE TABLE_NAME='table_name';
      ```

      where:

      - *retention_period* is the new retention period, in days
      - *table_name* is the name of the table you are changing

   b. Commit the changes to the database.

      Run the following from the **vsql** command line:

      ```
      COMMIT;
      ```

   For example, to change the data retention period of the RESOURCE_METRICS_BUILTIN table to 2192 days, run the following from the **vsql** command line:

   ```
   UPDATE SYS_TABLES_TO_PARTITION
   SET DATA_DAYS_RANGE='2192'
   WHERE TABLE_NAME='RESOURCE_METRICS_BUILTIN';
   COMMIT;
   ```

Splitting data transformer tasks to disperse workload

By default, there are two default scheduled tasks that control data transformers. Four scheduled tasks might not be enough to be able to run all the data transformers within one hour, so to enhance performance, you can split these data transformers into more tasks.

About this task

Table 10 shows a recommended format for splitting your data transformers into five tasks. The examples will make use of this table (specifically, with Task 1).

Table 10. Format for splitting data transformers into multiple tasks

Task	Data transformer name	Data flow entry
1	ClusterCapacity	`main_cluster_capacity.xml`
2	WorkloadStatistics	`main_workload_statistics.xml`
3	FlexLMLicusage	`main_flexlm_licusage.xml`
4	Hardware	`main_hardware.xml`
5	Jobmart	`main_jobmart.xml`

Procedure

1. Log in to the Analytics server host.
2. Create and enable a new scheduled task in the Platform Analytics Console:
   a. Launch the Platform Analytics Console.
      - Windows: **Start** > **All Programs** > **>IBM Corporation** > **Analytics Server** > **Analytics Console**
   b. Click **Scheduled Tasks** in the navigation tree.
   c. Right-click the main window and select **Add Scheduled Task**.
   d. Complete the required fields for the new task.
      - **Scheduled Task**: Specify the name of this task.
      - **Script File**: Specify `bin/dataagghourly.js` for hourly tasks or `bin/dataaggdaily.js` for daily tasks.
      - **Script Function**: Specify `doit`.

      For example, if you are creating Task 1 from the table with the recommended format of splitting data transformers, specify the following:
      - **Scheduled Task**: Specify **Task1** as the name of the scheduled task.
      - **Script File**: Specify `bin/dataagghourly.js` as the path to the script file.
      - **Script Function**: Specify `doit` as the script function.
   e. Enable the new scheduled task that you created.
3. In the `tasks` subdirectory of *`ANALYTICS_TOP`*, create a new directory with the same name as the name of the new scheduled task and navigate to the new directory.
4. From the new directory, create a text file of any name and with the `.tsk` extension.

   For example, create `task1.tsk`.
5. In the new `.tsk` text file, for each data transformer that you would like the scheduled task to control, add its corresponding data flow entry as a new file to the line.

You can also add a comment with the name of the data transformer if you start the line with the # character.

For example, for Task 1, the `task1.tsk` file should contain the following lines:

```
# Cluster Capacity
datatransformer/flow/clustercapacity/main_cluster_capacity.xml
```

Optimizing the Analytics node

You can perform the following optional tasks to enhance the performance of the Analytics node.

- "Increasing JVM memory"
- "Distributing the Analytics node workload"
- "Optimizing specific data loaders" on page 34

Increasing JVM memory

You can perform the following task to increase the amount of memory for the Java Virtual Machine (JVM).

About this task

If you experience a problem with the JVM running out of memory, you can increase the JVM memory to at least 2 GB for the Analytics node, instead of using the default value.

Note: Java cannot use more than 1638 MB of memory on 32-bit platforms. If you are using a 32-bit Windows system, you cannot increase the JVM memory beyond 1638 MB.

Procedure

1. Log in to the Analytics node host.
2. Edit the *ANALYTICS_TOP*`/conf/wsm/wsm_plc.conf` file.
3. Change the Java starting options to increase the JVM memory.

   Navigate to **JAVA_OPTS** and increase the JVM memory to at least 2 GB.

   For example:

   ```
   JAVA_OPTS=-Xms64m -Xmx2048m
   ```

Distributing the Analytics node workload

You can improve performance by distributing the Analytics node workload among multiple hosts or multiple loader controllers within a host so that each host or loader controller is responsible for a specific type (or types) of data loading.

About this task

Based on the type of data, you can categorize all the data loaders into different types, and enable each node or loader controller to be responsible for only one type of data loader. Table 11 on page 32 describes the different data loader categories and their corresponding loader controller configuration files.

Table 11. Data loader categories and corresponding loader controller configuration files

Data loader category	Loader controller configuration file
Host-related	• `plc_ego.xml` • `plc_coreutil.xml`
Job-related	• `plc_lsf.xml` • `plc_bjobs-sp012.xml`
Advanced job-related	• `plc_lsf_advanced.xml`
License-related	• `plc_license.xml`

You can perform the following tasks to distribute the Analytics node workload to multiple hosts or to multiple loader controllers in one node.

Procedure

Use one of the following procedures to distribute the Analytics node workload:

- "Distributing the Analytics node workload to multiple hosts"

  Distribute the workload to multiple hosts if your Analytics node is experiencing performance issues and cannot handle all the workload by itself.
- "Distributing the Analytics node workload to multiple loader controllers in one host" on page 33

  Distribute the workload to multiple loader controllers within a host to reduce performance bottlenecks if your Analytics node host is powerful enough to handle the workload (for example, the host has at least four cores and at least 6 GB of memory).

Distributing the Analytics node workload to multiple hosts

A single node might experience performance issues when handling high workload levels. To resolve this issue, you can distribute the node workload to multiple hosts so that each host is responsible for a specific type of data loading.

Procedure

1. Install and configure the Analytics node on other hosts in the same cluster.

   For example, if you plan to have one node for each data loader category, you should install and configure three additional Analytics node hosts. Therefore, your original node would handle host-related data loaders, while each of the other three nodes would handle each of the other three categories.

   Alternatively, you can have fewer nodes and still disperse workload to some degree. For example, you could have two nodes: your original node could handle host-related and job-related data loaders, while another node could handle advanced job-related and license-related data loaders.
2. For each Analytics node, disable the data loading categories that the node will not be handling:
   a. Navigate to the *ANALYTICS_TOP*`/conf/plc` directory.
   b. Move or rename the configuration files for the data loader categories which that Analytics node will no longer handle.

      For example, if you want the Analytics node to only handle host-related data loaders, move or rename all configuration files except `plc_ego.xml` and `plc_coreutil.xml`.
3. Restart the new loader controllers:

```
perfadmin stop all
perfadmin start all
```

Note: To stop or start an individual data loader, use `perfadmin stop` *plc_name* and `perfadmin start` *plc_name*.

For example, to stop the loader controller that handles license-related data loaders, run: `perfadmin stop plc_license`

Distributing the Analytics node workload to multiple loader controllers in one host

If your node host is powerful enough to handle the workload (for example, the host has at least four cores and at least 6 GB of memory), you can perform this task to reduce performance bottlenecks by creating multiple loader controllers and having each loader controller be responsible for a specific type of data loading. This enables each of the multiple cores in your host to control a single loader controller and be more efficient in sharing the workload among multiple cores.

Procedure

1. For each data loader category beyond the first, create a new loader controller on the same host.

   Since there are four data loader categories, create three additional loader controllers.

   In this task, *PLC_NAME* represents the name of the new loader controller that you will create.

   For example, you can create **plc_job** to handle job-related data loaders, **plc_advanced_job** to handle advanced job-related data loaders, and **plc_license** to handle license-related data loaders, while the original loader controller (**plc**) handles host-related data loaders.

   a. In the *ANALYTICS_TOP*`/conf/wsm` directory, copy the `wsm_plc.conf` file to a new file, one for each new loader controller.

      For each new loader controller, name the new file `wsm_`*PLC_NAME*`.conf`

      For example, copy `wsm_plc.conf` to `wsm_plc_job.conf`, `wsm_plc_advanced_job.conf`, and `wsm_plc_license.conf`.

   b. Edit each new `wsm_`*PLC_NAME*`.conf` file and specify the new loader controller name.

      Navigate to the **SERVICE_NAME**, **SERVICE_COMMAND**, and **LOG_PREFIX** parameters to use the new loader controller name as follows:

      ```
      SERVICE_NAME=PLC_NAME
      SERVICE_COMMAND=com.platform.perf.dataloader.Main –f PLC_NAME.xml
      LOG_PREFIX=PLC_NAME
      ```

      For example, for the `wsm_plc_job.conf` file:

      ```
      SERVICE_NAME=plc_job
      SERVICE_COMMAND=com.platform.perf.dataloader.Main –f plc_job.xml
      LOG_PREFIX=plc_job
      ```

      Make similar edits to the `wsm_plc_advanced_job.conf` and `wsm_plc_license.conf` files.

   c. In the *ANALYTICS_TOP*`/conf` directory, create a new subdirectory for each new loader controller.

      Name each subdirectory *PLC_NAME*.

      For example, create three new subdirectories named `plc_job`, `plc_advanced_job`, and `plc_license`.

d. Move the loader controller configuration files corresponding to the data loader categories from *ANALYTICS_TOP*/conf/plc to the new subdirectory corresponding to the new loader controller that will handle the category data loaders.

   For example:

   1) Move plc_lsf.xml and plc_bjobs-sp012.xml from *ANALYTICS_TOP*/conf/plc to *ANALYTICS_TOP*/conf/plc_job.
   2) Move plc_lsf_advanced.xml from *ANALYTICS_TOP*/conf/plc to *ANALYTICS_TOP*/conf/plc_advanced_job.
   3) Move plc_license.xml from *ANALYTICS_TOP*/conf/plc to *ANALYTICS_TOP*/conf/plc_license.

e. In the *ANALYTICS_TOP*/conf directory, copy the plc.xml loader controller file to a new file, one for each new loader controller.

   For each new loader controller, name the new file *PLC_NAME*.xml.

   For example, copy plc.xml to plc_job.xml, plc_advanced_job.xml, and plc_license.xml.

f. Edit each new loader controller file and change the **Port** parameter to a new port and change the **PLCDir** parameter to the new loader controller directory.

   For example:

   1) In plc_job.xml, change the **Port** value from 4046 to 4047, and change the **PLCDir** value from plc to plc_job.
   2) In plc_advanced_job.xml, change the **Port** value from 4046 to 4048, and change the **PLCDir** value from plc to plc_advanced_job.
   3) In plc_license.xml, change the **Port** value from 4046 to 4049, and change the **PLCDir** value from plc to plc_license.

2. Restart the new loader controllers:

```
perfadmin stop all
perfadmin start all
```

   Note: To stop or start an individual data loader, use perfadmin stop *plc_name* and perfadmin start *plc_name*.

   For example, to stop the loader controller that handles license-related data loaders, run: perfadmin stop plc_license

Optimizing specific data loaders

If you are encountering specific problems or are working under specific environments, you can perform this task to optimize certain data loaders to enhance the working performance of these individual data loaders.

Procedure

1. Optimize the FLEXnet usage data loader (**flexlicusageloader**) to improve data loading time if it cannot finish loading the data within one sampling interval.

   If the FLEXnet usage data loader cannot finish the data loading of one sampling interval within the time of one sampling interval (typically five minutes), the data loader will be behind. You need to enable the multi-threads setting to catch up with the multi-servers and daemons workload scale.

   a. Log into the Analytics node host that is running the FLEXnet usage data loader.
   b. Edit the *ANALYTICS_TOP*/conf/dataloader/flexlicusage.properties file.

c. Edit the **ThreadNumber** parameter to increase the number of threads for collecting data concurrently.

   For example:

   ```
   ThreadNumber=10
   ```

2. Optimize the FLEXnet Manager data loader (**fnmloader**) to achieve a larger capacity if you need to handle more than 1 million events per day.

   Using the default configuration means that you do not have to modify the data schema of the FLEXnet Manager reporting database, but the capacity of this data loader is limited to 1 million events per day.

   If you need to achieve a larger capacity, you need to modify the data schema of the FLEXnet Manager reporting database to increase performance. For details, see the *ANALYTICS_TOP*/conf/dataloader/FNMLOADER.readme file.

3. Optimize the LSF events data loader (**lsfeventsloader**) to prevent data loss if you are using LSF version 7.0.3 or earlier.

   The LSF events data loader reads data from the lsb.stream file. If you are using LSF version 7.0.3 or earlier, increase the size of the lsb.stream file to 2 GB or more to prevent data loss.

   a. Log in to the Analytics node host that is running the LSF events data loader.

   b. Edit the *LSF_ENVDIR*/lsbatch/*cluster_name*/configdir/lsb.params file.

   c. In the **Parameters** section, define a new parameter named **MAX_EVENT_STREAM_SIZE**.

      For example:

      ```
      Begin Parameters
      ⋮
      MAX_EVENT_STREAM_SIZE = 2048
      ⋮
      End Parameters
      ```

   d. Reconfigure the LSF cluster for your changes to take effect:

      ```
      lsadmin reconfig
      ```

Optimizing the Analytics database

You can perform this task to optimize the operation of the Analytics database.

Procedure

1. Increase the maximum number of client sessions.

   You might need to increase the maximum number of client sessions for the database to communicate with the Platform Analytics nodes. In Vertica, this is the **SESSIONS** parameter, which you can change from the **vsql** command line by using the **SET_CONFIG_PARAMETER** function:

   ```
   SELECT SET_CONFIG_PARAMETER('MaxClientSessions', SESSIONS_value);
   ```

   Each Analytics node needs approximately 30 sessions and the Analytics server needs approximately 10 sessions. In addition, each intended user that will concurrently access the live Platform Analytics reports will also need a client session.

   Therefore, for optimal performance of the database, calculate the optimal value of the **SESSIONS** parameter in Analytics as:

   SESSIONS_value = (*# of Analytics nodes*) × 30 + 10 + (*# of concurrent users browsing live reports*)

   For example, if you have three Analytics nodes, and will have an extra ten users concurrently accessing live Platform Analytics reports, increasing the

maximum number of client sessions to 110 should be sufficient. From the **vsql** command line, run the following command:

```
SELECT SET_CONFIG_PARAMETER('MaxClientSessions', 110);
```

2. Modify the default queue timeout of the **tm** resource pool to 90 seconds.

   From the **vsql** command line, run the following command:

   ```
   ALTER RESOURCE POOL tm queuetimeout 90;
   ```

3. To obtain optimal performance, store TEMP and DATA in different storage locations.

   TEMP data is distributed across available storage locations based on available storage space. However, DATA can be stored on different storage locations based on predicted or measured access patterns.

4. For each database host in the cluster, set the block device (**blockdev**) size.

   The block device is the physical storage device on the database. Set the **blockdev** size to obtain optimal performance by running the following commands for the drives on which your data directory is located:

   ```
   sudo blockdev --getss drive
   sudo blockdev --getra drive
   ```

   For example:

   ```
   sudo blockdev --getss /dev/md0
   sudo blockdev --getra /dev/md0
   ```

   By default, the first command should return 512. The second command should return 2048, which means the **readahead** parameter is set to 1 MB.

   If the **readahead** parameter is set too high, the database host may experience a slow mergeout.

5. *(Optional)* If the **readahead** parameter is set too high, set the **blockdev** size to correct this problem:

   ```
   sudo blockdev --setra 2048
   ```

   To retain this setting every time the host restarts, copy this line into the /etc/rc.local file.

Chapter 5. Uninstalling Platform Analytics

Perform this task if you want to uninstall Platform Analytics from a host. If you want to uninstall the Analytics database, refer to the *Installation Guide* for the Vertica Analytic Database.

Uninstalling the Analytics server

Perform this task to uninstall the Analytics server from a host.

About this task

This will remove all previously installed Platform Analytics services, reports, and programs from the host.

Procedure

1. Log in to the Analytics server host with the same account that you used to install the Analytics server.
2. Navigate to the *ANALYTICS_TOP*\uninst directory.
3. Run the uninstall script:

   ```
   uninstaller.bat
   ```

Chapter 6. Troubleshooting Platform Analytics

The troubleshooting procedures help you to resolve a variety of common problem scenarios.

Some or all Analytics database nodes fail to start due to a memory error

Some or all Analytics database nodes may fail to start and issue the following error:

```
Large:Memory(KB) Exceeded: Requested = number, Free = number
```

This error occurs because of a problem with the Resource Manager in the Analytics database. To resolve the problem, disable the Resource Manager before running the database, then enable the Resource Manager after the database has started. The resolution method depends on whether all database nodes failed to start, or only some database nodes failed to start.

Starting the Analytics database if all nodes failed to start

Use this procedure to start the database in the case where all of the nodes have failed to start.

Procedure

1. Manually disable the Resource Manager on all Analytics database nodes.

   Perform the following steps on each host in the database cluster:

   a. Log in to a host in the database cluster.

   b. Navigate to the directory containing the `vertica.conf` file.

      The directory is the directory containing the catalogs of the database that you want to start. This is the **Catalog pathname** that you were initially prompted to specify when you first created the database.

   c. Edit the `vertica.conf` file and add the following line to the end of the file:

      ```
      EnableResourceManager=0
      ```

2. Start the Analytics database on all database nodes.
3. Re-enable the Resource Manager:

   a. Log in to a host in the database cluster.

   b. Run the following SQL statement from the **vsql** command line:

      ```
      SELECT SET_CONFIG_PARAMETER('EnableResourceManager', '1');
      ```

Starting the Analytics database if some nodes failed to start

Use this procedure to start the database in the case where only some of the nodes have failed to start.

Procedure

1. Log in to an Analytics database node that is still running.
2. Disable the Resource Manager.

   Run the following SQL statement from the **vsql** command line:

   ```
   SELECT SET_CONFIG_PARAMETER('EnableResourceManager', '0');
   ```

3. Start the Analytics database on all of the database nodes that failed to start.
4. Re-enable the Resource Manager:
   a. Log in to a host in the database cluster.
   b. Run the following SQL statement from the **vsql** command line:
      ```
      SELECT SET_CONFIG_PARAMETER('EnableResourceManager', '1');
      ```

Analytics node does not send events after installation if started before the Analytics server

After installing Platform Analytics using a clean database, if you start the Analytics node before starting the Analytics server, the node will not send events. This problem only occurs the first time after installation.

The `EVENT_MANAGER_CONF` table for the event locator is not initialized until you start the Analytics server for the first time. Therefore, if you start the Analytics node without first starting the Analytics server after initial installation with a clean database, the event sender does not have access to the `EVENT_MANAGER_CONF` TABLE until you start the Analytics server.

To resolve this problem, restart the Analytics node after you start the Analytics server.

FLEXnet usage data loader cannot obtain license usage data due to insufficient swap space

If you have an Analytics node running on a UNIX host, the FLEXnet usage data loader (**flexlicusageloader**) log may report `Failed to obtain license usage from the license server` and `Not enough space` errors. This problem does not apply to Windows hosts.

This error occurs if you have insufficient disk space allocated to the swap space on that host. To work around this issue, extend the swap space so that it has at least 2 GB of free space on that host before starting the Analytics node on the host.

In certain configurations, the Platform Analytics Console shows that the loader controller is down, but perfadmin shows it is running

In the Platform Analytics Console, if you click **Data Collection Nodes**, you may see that the loader controller is `Down`. However, if you examine the loader controller service (**plc**) in the Platform Analytics node (using **perfadmin list**) the loader controller service is `STARTED`.

This can occur if the loopback IP address (`127.0.0.1`) is incorrectly defined using the name of your host rather than `localhost` in the `/etc/hosts` file, or if your host has multiple network interface cards (NICs).

To resolve this problem, change the loopback IP address and NSS (Name Service Switch) configuration.

Changing the loopback IP address and NSS configuration

Use this procedure to change the loopback IP address and NSS configuration.

Procedure

1. Change the loopback IP address to `localhost`:
   a. Edit the `/etc/hosts` file.
   b. Navigate to the line with the definition for the loopback IP address (`127.0.0.1`).

      If this IP address is not defined as `localhost`, you must change the definition.

      For example, if your host is `hostA` in the `example.com` domain, you must navigate to and change the following line:

      ```
      127.0.0.1       hostA       hostA.example.com
      ```

   c. Either delete the line or change the definition to `localhost` and save the file.

      For example, either delete the line or change it to the following:

      ```
      127.0.0.1       localhost   localhost.localdomain
      ```

2. If your host has multiple network interface cards (NICs), change the NSS (Name Service Switch) configuration to look up NIS before looking up the file for host names and numbers:
   a. Edit the `/etc/nsswitch.conf` file.
   b. Navigate to the line with the definition for `hosts`.

      For example, by default, this line is:

      ```
      hosts:      files nis dns
      ```

   c. Change the line so `nis` appears before `files` and save the file.

      For example:

      ```
      hosts:      nis files dns
      ```

3. Restart the services on the Analytics node.

   ```
   perfadmin stop all
   perfadmin start all
   ```

Failure during Platform Analytics installation

A crashed InstallShield database may cause the Platform Analytics installation to fail. If the Platform Analytics installation failed to complete, you may need to manually remove the InstallShield Multi-Platform (ISMP) database.

Remove the ISMP database from the following directories:

- Windows: `C:\Program Files\Common Files\InstallShield\Universal\common`
- UNIX: `~/InstallShield`

Cluster Capacity and Workload Statistics workbook displays only the first execution host in the execution host list for parallel jobs

This is applicable only for version 7.*x* clusters.

The Cluster Capacity and Workload Statistics workbook displays parallel job execution hosts as one host and gets the data from the first execution host, even though parallel jobs are running on different hosts. For example, if a parallel job execution host is `3*hostA 4*hostB`, the cluster capacity data transformer assumes that 7 slots are occupied by host A.

Incorrect number of down slots reported

If the number of job slots is defined using "!" in the **lsb.hosts** file and a host is down, then the number of down slots reported is not correct. To work around this issue, define the number of slots for each host in the cluster in **lsb.hosts**.

Inaccurate license usage data collected for the license vendor daemons

If you have multiple license vendor daemons on a license server sharing the same port, the license usage data for those license vendor daemons may not be correct. To work around this issue, download the older version of **lmutil** from the Platform FTP site.

Download the **lmutil** binary from `patches/lsf_analytics/8.0/FLEXlm9.2/<platform>` and move it to `ANALYTICS_TOP/license/7.0/<platform>/...`

Cannot install the Analytics node

Installation of the Analytics node will fail when the **LSF_VERSION** defined in the `lsf.conf` file is not the actual version.

To resolve this problem, before you install the Analytics node, edit the `lsf.conf` file and change the value of **LSF_VERSION** to an appropriate version. For example, if the actual LSF version is 7.*x* but the **LSF_VERSION** in `lsf.conf` is set to `active`, then before you install the node, change the **LSF_VERSION** to `7.0`. After installing the node, change the **LSF_VERSION** back to `active`.

Troubleshooting third-party issues

Use this information to troubleshoot some common third-party problem scenarios.

Message "Out of memory" displays after clicking the Data tab

This error message is displayed when you try to view big data, which is more than 4 GB in size. To avoid the error, you can either narrow down the data range or increase the memory size of the host.

Average data in the Cluster Usage table is less accurate than data on the Cluster Usage graph

Data shown in the table is not accurate in some of the roll-up levels, as it considers the sampling points of data instead of the whole date period.

For example, Table 12 shows sampling points of data for slots number with different slot status:

Table 12. Example of sampling data points for slots number with different slot status

Sampling points for different slot status	10:00	10:10	10:20	10:30	10:40	10:50
RUN					1	1
DOWN	2	2	2	2	2	2

In the Cluster Usage table, the average slots number for the RUN status, rolled up for hour 10, is:

```
(1 + 1) / 2 = 1
```

However, the graph data shows the correct value, which is:

```
(1 + 1) / 6 = 0.33
```

As a workaround, refer to the Cluster Usage graph for more accurate data.

The Projects dashboard of the Workload Accounting report reports session busy error when sorting big data

In the Projects dashboard of the Workload Accounting report, if you select big data and try to sort, the reporting server may display the following error:

```
Unexpected Server Error: Session busy, please try later.
```

For example, if you select data for more than 3 years and try to drill down to a specific year that has more than 20 K projects and sort as project name, you will see that error.

To avoid the error, narrow down the data range or try to view using the Platform Analytics Designer.

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.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web

sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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
Intellectual Property Law
Mail Station P300
2455 South Road,
Poughkeepsie, NY 12601-5400
USA

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 document 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_.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com® are 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 http://www.ibm.com/legal/copytrade.shtml.

LSF, Platform, and Platform Computing are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a 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.

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other company, product, or service names may be trademarks or service marks of others.

Printed in USA

SC14-7571-00