
Administering Platform Analytics

Platform Analytics
Version 8.0.2
Release date: September 2011



Copyright

© 1994-2011 Platform Computing Corporation.

Although the information in this document has been carefully reviewed, Platform Computing Corporation (“Platform”) does not warrant it to be free of errors or omissions. Platform reserves the right to make corrections, updates, revisions or changes to the information in this document.

UNLESS OTHERWISE EXPRESSLY STATED BY PLATFORM, THE PROGRAM DESCRIBED IN THIS DOCUMENT IS PROVIDED “AS IS” AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL PLATFORM COMPUTING BE LIABLE TO ANYONE FOR SPECIAL, COLLATERAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, INCLUDING WITHOUT LIMITATION ANY LOST PROFITS, DATA, OR SAVINGS, ARISING OUT OF THE USE OF OR INABILITY TO USE THIS PROGRAM.

We’d like to hear from you

You can help us make this document better by telling us what you think of the content, organization, and usefulness of the information. If you find an error, or just want to make a suggestion for improving this document, please address your comments to doc@platform.com.

Your comments should pertain only to Platform documentation. For product support, contact support@platform.com.

Document redistribution and translation

This document is protected by copyright and you may not redistribute or translate it into another language, in part or in whole.

Internal redistribution

You may only redistribute this document internally within your organization (for example, on an intranet) provided that you continue to check the Platform Web site for updates and update your version of the documentation. You may not make it available to your organization over the Internet.

Trademarks

LSF is a registered trademark of Platform Computing Corporation in the United States and in other jurisdictions.

ACCELERATING INTELLIGENCE, PLATFORM COMPUTING, PLATFORM SYMPHONY, PLATFORM JOB SCHEDULER, PLATFORM ISF, PLATFORM ENTERPRISE GRID ORCHESTRATOR, PLATFORM EGO, and the PLATFORM and PLATFORM LSF logos are trademarks of Platform Computing Corporation in the United States and in other jurisdictions.

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

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Microsoft is either a registered trademark or a trademark of Microsoft Corporation in the United States and/or other countries.

Windows is a registered trademark of Microsoft Corporation in the United States and other countries.

Intel, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other products or services mentioned in this document are identified by the trademarks or service marks of their respective owners.

Third-party license agreements

<http://www.platform.com/Company/third.part.license.htm>

Contents

Part I: About Platform Analytics ... 5

1	Introduction to Platform Analytics	7
2	Architecture overview	9
3	System architecture	13
4	PERF directories in the Platform Analytics node	15
5	Licensing	17

Part II: Managing database host ... 19

6	Database	21
7	Data sources	23

Part III: Managing Platform Analytics node ... 27

8	Loader controller	29
9	Data loaders	33
10	Platform Analytics node command-line tools	41
11	Platform Analytics node configuration files	45

Part IV: Managing Platform Analytics server ... 51

12	Platform Analytics Console	53
13	Data transformers	55
14	Event notification	59
15	Data purger	63
16	Scheduled tasks	67
17	Platform Analytics server command-line tools	73
18	Platform Analytics server configuration files	77

Part V: Viewing reports ... 85

19	Generating reports	87
20	Collecting data and viewing reports	89

21	Platform Application Center (optional)	93
----	--	----

Part VI: Managing Platform Analytics ... 101

22	Secure your data and working environment	103
23	Maintaining the database	105
24	Troubleshooting the node	109
25	Troubleshooting the server	117

Part VII: Customizing Platform Analytics ... 121

	Naming conventions	123
	Node customizations	124
	Server customizations	127
	Database schema customizations	128
	Customization management	129



About Platform Analytics

Platform Analytics provides several interactive dashboards that are ready to use “out of the box”, making it quick and easy to analyze key data. Existing or new data sources can be rapidly combined with Platform Analytics data to provide data views tailored specifically to an organization’s unique requirements without the need to build intermediate data views.

Introduction to Platform Analytics

Platform Analytics is an advanced analysis and visualization tool for analyzing massive amounts of Platform LSF workload data. It enables managers, planners and administrators to easily correlate job, resource and license data from one or multiple Platform LSF clusters for data-driven decision-making. With better insight into HPC datacenter environment, organizations can identify and quickly remove bottlenecks, spot emerging trends and plan capacity more effectively.

Unlike traditional business intelligence solutions that require significant time and multiple steps to translate raw data into usable information, Platform Analytics incorporates innovative visualization tools that are built on top of a powerful analytics engine for quick and easy results. Users can utilize the pre-configured dashboards or construct their own, quickly answer questions about their HPC infrastructure and applications and use that information to optimize HPC resource utilization.


Platform Analytics is a workload intelligence solution for LSF cluster, FLEXnet license, and FLEXnet Manager license data. Platform Analytics collects LSF and license data, then assembles it into reports for your analysis. Platform Analytics provides all the tools you need to collect the data, load it into a database, then convert it to reports for your analysis using a ROLAP (relational online analytical processing) tool.

Architecture overview

The Platform Analytics architecture is based on the Platform Enterprise Reporting Framework (PERF) architecture. Platform Analytics adopts and extends the PERF technology to cover all data collection requirements and to improve data collection reliability. Platform Analytics supports Vertica, a state-of-the-art MPP columnar database that runs on standard hardware and uses a fraction of the resources of traditional database management systems. Support for the Oracle database is also available. The Platform Analytics reporting server that has Tableau Server is used as the Relational Online Analytics Processing (ROLAP) tool to generate reports and to allow other users to view these reports using a web browser.

Major Components of Platform Analytics

The following are the components of Platform Analytics.



Platform Analytics Data Loader

A Data Loader is installed on each cluster. This loader helps to load data directly into Database. Each data loader collects LSF data, FlexLM License Data (from any number of Flex LM License Servers), and FNM License Data (from a FNM License Server).

Database - Vertica / Oracle

Platform Analytics is designed to support the Vertica database, to provide improvements in query and data loading performance over traditional RDBMS technologies. Data is neatly organized into tables for reporting and analysis.

Platform Analytics also supports Oracle database for data storage and analysis.

Platform Analytics Server

Platform Analytics Server communicates between the data loaders and the Vertica database. It manages the data which the Platform Analytics nodes collect. The Platform Analytics Server receives event notification from nodes and other components, and then sends out an email according to the rule configured.

Platform Analytics Reporting Server

Platform Analytics Reporting Server is a web based reporting tool consisting of Workbooks. It collects data from the Vertica database and allows the publishing of Dashboards or individual Worksheets from the Platform Analytics Designer.

Platform Analytics Designer (Optional)


Platform Analytics Designer provides users with the flexibility to easily construct complex queries and dashboards specific to each customer's own reporting and analysis requirements. This designer is mainly used for customizing existing Analytics workbooks and for creating new custom workbooks based on Analytics Database data.

Platform Application Center (Optional)

Platform Application Center is used to view the Platform Analytics reports. It allows you to look at the overall statistics of the entire cluster. Platform Application Center helps to analyze the history of hosts, resources, and workload in the cluster to get an overall picture of cluster's performance.

System architecture

The system architecture gives an overview of data flow in Platform Analytics.



System ports

For a list of ports that the Platform Analytics hosts use, refer to *Installing ; Platform Analytics* (specifically, the *System ports* section in the *Platform Analytics hosts* chapter).

PERF directories in the Platform Analytics node

PERF components reside in various `perf` subdirectories within the LSF directory structure. This document uses `LSF_TOP` to refer to the top-level LSF installation directory, and `ANALYTICS_TOP` to refer to the top-level Platform Analytics installation directory. In UNIX, you need to source the PERF environment to use these environment variables.

PERF directory environment variables in UNIX

Directory name	Directory description	Default file path
<code>\$PERF_TOP</code>	PERF directory	<code>ANALYTICS_TOP</code>
<code>\$PERF_CONFDIR</code>	Configuration files	<code>ANALYTICS_TOP/conf</code>
<code>\$PERF_LOGDIR</code>	Log files	<code>ANALYTICS_TOP/log</code>
<code>\$PERF_WORKDIR</code>	Working directory	<code>ANALYTICS_TOP/work</code>

PERF directory environment variables in Windows

Directory name	Directory description	Default file path
<code>%PERF_TOP%</code>	PERF directory	<code>ANALYTICS_TOP</code>
<code>%PERF_CONFDIR%</code>	Configuration files	<code>ANALYTICS_TOP\conf</code>
<code>%PERF_LOGDIR%</code>	Log files	<code>ANALYTICS_TOP\log</code>
<code>%PERF_WORKDIR%</code>	Working directory	<code>ANALYTICS_TOP\work</code>

PERF directories in the Platform Analytics node


Licensing

The Platform Analytics license file includes licenses for data collection (data volume audit for Vertica.)

Contact Platform Computing to obtain a license for Platform Analytics. You may purchase and enable the following components for your Platform Analytics installation to be included in the Platform Analytics license file:

Analytics base	This is a must have license for Platform Analytics. This license allows you to collect data from LSF clusters.
LSF advanced data collection	This license allows you to collect LSF advanced data from LSF clusters. LSF advanced data is cluster performance and operation data that is not gathered in the base PERF package included with LSF.
License data collection	This license allows you to collect license usage and event data from your FLEXnet servers.
Vertica database connector	This license allows you to monitor data volume for Vertica only.

If you have a demo license and obtained a production license, you need to replace the old demo license file in the node (*PERF_CONFDIR/license.dat*) and the server (*PA_SERVER_ROOT/conf/license.dat*) with the new demo license file. Make sure that the replaced file name is *license.dat*.



Managing database host

Managing database host

6

Database

The relational database contains the cluster data, organised into tables, for reporting and analysis.

About the database

The relational database contains the cluster operations data for reporting and analysis. Platform Analytics components input and output data from the tables within the database. Apart from Vertica, Platform Analytics supports Oracle 9i, 10g, and 11g databases.

Default behavior

Data is stored and organized in tables within the database. The organization of this data is defined in the data schema of the tables.

The database and its data schema are partitioned for Platform Analytics data. A partitioned database has tables divided into multiple, smaller tables. This improves database performance for larger clusters.


In a large database, purging old job records, transforming data, and other database maintenance tasks can have a significant effect on database performance. Purging old job records and transforming data from smaller tables has less of an impact on the system performance of active tables than on larger tables.

The database tables are partitioned by quarter. Platform Analytics keeps three years of data in the database. Every month, Platform Analytics has a scheduled task that drops any partition that is older than three years by quarter.

Database interactions

All interactions between Platform Analytics and the database are through the JDBC connection as defined by the data sources.

The following diagram illustrates the interaction between the database and other components.



Data sources

Data sources define the JDBC connections between the hosts and the database.

About data sources

Data sources define all JDBC connections to the data tables in the relational database. The data tables contain processed cluster data that can be extracted and used in reports.

You define the JDBC connection to the database when you install Platform Analytics. The information about the JDBC driver together with the user and password information is called the data source. If you change your database or modify your connection, you need to update the data source properties in Platform Analytics accordingly. The default Platform Analytics data source for the server and the node is ReportDB.

Platform Analytics uses one or more data sources. You must install JDBC drivers for your database type on the Platform Analytics server host before defining the corresponding data source.

Data source interactions

The data source is the JDBC connection between the data tables in the relational database and all Platform Analytics components. Any interaction with the data tables in the database goes through the JDBC connection as defined in the data source.

Server data source interactions

Data transformers obtain data from the data tables through the server data sources, and stores transformed data into the data tables through the server data sources.


The data purger purges old records from the data tables through the server data sources.

Node data source interactions

The data sources for the Platform Analytics node interact with the data tables in the database. If your cluster has multiple FLEXnet Manager servers, each FLEXnet Manager server has its own data source.

Data loaders either request cluster operation data, or obtain it directly from the data tables through the node data sources. The data loaders store this data into data tables through the node data sources.

The following diagram illustrates the interaction between data sources and other components.



Data source actions


Actions on the Platform Analytics server data sources

If the Platform Analytics server is running on a UNIX host, you need to restart the Platform Analytics server daemons (by running `perfadmi n stop all` and `perfadmi n start all`) after changing the server data source.

Action	Platform Analytics Console
View the list of server data sources.	In the navigation tree, click Data Sources .
Add a server data source.	When viewing the list of data sources, select Action > Add Data Source .
Edit the settings of a server data source.	When viewing the list of data sources, click the data source and select Action > Edit Data Source .
Delete a server data source.	When viewing the list of data sources, click the data source and select Action > Remove Data Source .

Actions on the Platform Analytics node data sources

Action	Command line
Add a node data source.	UNIX: dbconfig.sh add <i>data_source_name</i> Windows: dbconfig add <i>data_source_name</i> where <ul style="list-style-type: none"><i>data_source_name</i> is the name the data source that you want to add.
Edit the settings of the Platform Analytics node data source (ReportDB).	UNIX: dbconfig.sh Windows: dbconfig
Edit the settings of any node data source, including FLEXnet Manager data sources.	UNIX: dbconfig.sh edit <i>data_source_name</i> Windows: dbconfig edit <i>data_source_name</i> where <ul style="list-style-type: none"><i>data_source_name</i> is the name the data source that you want to edit.



Managing Platform Analytics node

Platform Analytics nodes are hosts that collect data from clusters or license servers. Each node either belongs to a cluster from which Platform Analytics collects data, or is a standalone host that collects license data.

8

Loader controller

The Platform loader controller (pl c) controls the data loaders that gather data from the system and writes the data into the relational database containing raw data.

About the loader controller

The loader controller manages the data loaders by controlling the schedule in which each data loader gathers data.

Logging levels

There are logging levels that determine the detail of messages that the PERF services record in the log files. In decreasing level of detail, these levels are ALL (all messages), TRACE, DEBUG, INFO, WARN, ERROR, FATAL, and OFF (no messages).

By default, the PERF services log messages of INFO level or higher (that is, all INFO, WARN, ERROR, and FATAL messages).

The loader controller log file is located in the log directory:

- UNIX: \$PERF_LOGDIR
- Windows: %PERF_LOGDIR%


Default behavior

The loader controller service starts automatically when the master host starts up if you have the loader controller registered as an RC.

Loader controller interactions

The loader controller service controls the scheduling of the data loaders. Sampling and retrieving data loaders request cluster operation data from the data tables through the node data sources while other data loaders obtain it directly from the data tables through the node data sources. The data loaders store this data into data tables through the node data sources. Each data loader contains data that is stored in specific data tables in the database.

The following diagram illustrates the interaction between the loader controller and other components.



Configuration to modify loader controller behavior

Action	Configuration files	Parameter and syntax
Specify the default log level of your <code>plc</code> log file.	<code>log4j.properties</code> File location: UNIX: <code>\$PERF_CONFDIR</code> Windows: <code>%PERF_CONFDIR%</code>	<code>log4j.logger.com.platform.perf.data.loader=log_level, com.platform.perf.data.loader</code> where <ul style="list-style-type: none"> <code>log_level</code> is the default log level of your loader controller log files. The loader controller only logs messages of the same or lower level of detail as <code>log_level</code> . Therefore, if you change the log level to <code>ERROR</code> , the loader controller will only log <code>ERROR</code> and <code>FATAL</code> messages.

Loader controller actions

Actions on the loader controller service

Note:

To stop or start the `plc` service, you must run the commands on the local host running the `plc` service.

Action	Command line
View the status of the <code>plc</code> and other <code>PERF</code> services.	<code>perfadmin list</code>
Stop the <code>plc</code> service.	<code>perfadmin stop plc</code>
Start the <code>plc</code> service.	<code>perfadmin start plc</code>

Actions to change the loader controller settings

Action	Command line
Dynamically change the log level of your loader controller log file (temporarily).	UNIX: <code>plcclient.sh -l log_level</code> Windows: <code>plcclient -l log_level</code> where <ul style="list-style-type: none"> <code>log_level</code> is the log level of your loader controller log file. If you restart the loader controller, these settings will revert back to the default level.
	<h4>Note:</h4> <p>You must run this command on the local host running the <code>plc</code> service.</p>

Loader controller

9

Data loaders

Data loaders gather cluster operation data and load it into tables in a relational database containing raw data. Data loaders are controlled by the Platform loader controller (pl c) service.

About data loaders

Data loaders are polling loaders or history data loaders. The data loaders gather data and load this data into specific tables in the relational database containing raw data. Data loaders handle daylight savings automatically by using GMT time when gathering data.

Logging levels

There are logging levels that determine the detail of messages that the data loaders record in the log files. In decreasing level of detail, these levels are ALL (all messages), TRACE, DEBUG, INFO, WARN, ERROR, FATAL, and OFF (no messages).

By default, data loaders log messages of INFO level or higher (that is, all INFO, WARN, ERROR, and FATAL messages).

The data loader log files are located in the `data loader` subdirectory of the log directory:

- UNIX: `$PERF_LOGDIR/data loader`
- Windows: `%PERF_LOGDIR%\data loader`

Default behavior

Data loaders gather data from data sources at regular intervals. The following are lists of the data loaders, the specific loader controller configuration file (`plc_*.xml`), and the default behavior:

LSF host data loaders (`plc_coreutil.xml`)

Data loader name	Data type	Data gathering interval	Data loads to	Loader type
Host core utilization (<code>hostcoreutil loader</code>)	core utilization	5 minutes	HOST_CORE_UTILIZATION	polling

LSF job data loaders (`plc_bjobs-sp012.xml`)

Data loader name	Data type	Data gathering interval	Data loads to	Loader type
Bjobs (<code>lsfbjobs loader</code>)	job-related	10 minutes	LSF_BJOBS	polling

LSF data loaders (`plc_lsf.xml`)

Data loader name	Data type	Data gathering interval	Data loads to	Loader type
Host metrics (<code>hostmetrics loader</code>)	host-related metrics	10 minutes	RESOURCE_METRICS RESOURCES_RESOURCE_METRICS	polling

Data loader name	Data type	Data gathering interval	Data loads to	Loader type
Host properties (hostpropertiesloader)	resource properties	1 hour	LSF_RESOURCE_PROPERTIES	polling
Bhosts (lsfbhostsl oader)	host utilization and state-related	10 minutes	LSF_BHOSTS	polling
LSF events (lsfeventsl oader)	events with a job ID, performance events, resource events, JOB_FINISH2 events	5 minutes	LSB_EVENTS LSB_EVENTS_EXECHOSTLIST LSF_PERFORMANCE_METRIC LSB_JOB_FINISH & LSB_JOB_EXECHOSTS	file
Resource properties (lsfrespropl oader)	shared resource properties	1 hour	LSF_RESOURCE_PROPERTIES	polling
SLA (lsfslal oader)	SLA performance	5 minutes	LSF_SLA	polling
Shared resource usage (sharedresusage l oader)	shared resource usage	5 minutes	SHARED_RESOURCE_USAGE SHARED_RESOURCE_USAGE_HOSTLIST	polling

LSF advanced data loaders (plc_lsf_advanced.xml)

Data loader name	Data type	Data gathering interval	Data loads to	Loader type
Host group (hostgroup l oader)	host group	1 hour	HOST_GROUP	polling
Bqueues (lsfbqueue l oader)	queue properties	5 minutes	LSF_BQUEUES	polling
Pending reason (lsfpendi ngreason l oader)	job pending reasons	15 minutes	JOBS_PENDING_REASON DPR_BYINTERVAL	polling
User group (usergroup l oader)	user group	1 hour	USER_GROUP	polling
Pending Reasons (lsbpendi ngreasonsl oader)	job pending reason - from the LSF data file lsb.pendi ngreas ons	10 minutes	LSB_JOB_PENDINGREASON	file
Job status (lsfjobst atus l oader)	job status - from the LSF data file lsb.st atus	10 minutes	LSB_JOB_STATUS	file


FLEXnet data loaders (pl c_ l i cense. xml)

Data loader name	Data type	Data gathering interval	Data loads to	Loader type
FLEXnet usage (f l e x l i c u s a g e l o a d e r)	license usage	5 minutes	FLEXLM_LICENSE_USAGE	polling
FLEXnet events (f l e x l i c e v e n t s l o a d e r)	license log file event	5 minutes	FLEXLM_LICENSE_EVENTS	file
FLEXnet Manager (f n m l o a d e r) - only supports FLEXnet Manager 11 or later.	license event	30 minutes	FLEXNET_LICENSE_EVENTS	database

Data loader interactions

The loader controller service controls the scheduling of the data loaders. The data loaders store LSF and license data into data tables through the node data sources. Each data loader contains data that is stored in specific data tables in the database.

The following diagram illustrates the interaction between data loaders and other components.



Configuration to modify data loader behavior

After editing the loader controller configuration files, restart the loader controller for your changes to take effect. The specific loader controller configuration file (pl c_*. xml) depends on the type of data loader.

These files are located in the loader controller configuration directory:

- UNIX: \$PERF_CONFDIR/pl c
- Windows: %PERF_CONFDIR%\pl c

Action	Configuration files	Parameter and syntax
Specify the frequency of data gathering for the specified data loader.	Loader controller configuration files for your data loaders (pl c_*. xml).	<pre><DataLoader Name="loader_name" Interval="gather_interval" ... /></pre> <p>where</p> <ul style="list-style-type: none"> • <i>loader_name</i> is the name of your data loader • <i>gather_interval</i> is the time interval between data gathering, in seconds
Enable data gathering for the specified data loader. This is enabled by default.		<pre><DataLoader Name="loader_name" ... Enable="true" ... /></pre> <p>where</p> <ul style="list-style-type: none"> • <i>loader_name</i> is the name of your data loader
Disable data gathering for the specified data loader.		<pre><DataLoader Name="loader_name" ... Enable="false" ... /></pre> <p>where</p> <ul style="list-style-type: none"> • <i>loader_name</i> is the name of your data loader
Enable data loss protection for the specified data loader. This is enabled by default.	Specific data loader configuration file: <i>data_loader_name.xml</i>	<pre><Writer ... EnableRecover="Y"></pre>
Disable data loss protection for the specified data loader.	File location: UNIX: \$PERF_CONFDIR/ data_loader Windows: %PERF_CONFDIR %\data_loader	<pre><Writer ... EnableRecover="N"></pre>

Action	Configuration files	Parameter and syntax
Specify the default log level of your data loader log files.	<p><code>log4j.properties</code></p> <p>File location:</p> <p>UNIX: <code>\$PERF_CONFDIR</code> Windows: <code>%PERF_CONFDIR%</code></p>	<pre>log4j.logger.\${data loader}=log_level, \${data loader}</pre> <p>where</p> <ul style="list-style-type: none"> <code>log_level</code> is the default log level of your data loader log files.
Specify the log level of the log files for the specified data loader.		<pre>log4j.logger.data loader.loader_name=log_level</pre> <p>where</p> <ul style="list-style-type: none"> <code>loader_name</code> is the name of the data loader. <code>log_level</code> is the log level of the specified data loader. <p>For example, to set the LSF events data loader (<code>lsfeventsloader</code>) to ERROR, add the following line to <code>log4j.properties</code>:</p> <pre>log4j.logger.data loader.lsfeventsloader=ERROR</pre>
Specify the log level of the log files for the reader or writer area of the specified data loader.		<pre>log4j.logger.data loader.loader_name.area=log_level</pre> <p>where</p> <ul style="list-style-type: none"> <code>loader_name</code> is the name of the data loader. <code>area</code> is either reader or writer. <code>log_level</code> is the log level of the specified data loader. <p>For example, to set the LSF events data loader (<code>lsfeventsloader</code>) writer to DEBUG, add the following line to <code>log4j.properties</code>:</p> <pre>log4j.logger.data loader.lsfeventsloader.writer=ERROR</pre>

The data loaders only log messages of the same or lower level of detail as *log_level*. Therefore, if you change the log level to ERROR, the data loaders will only log ERROR and FATAL messages.

Data loader actions

Action	Command line
View the status and logging levels of the data loaders.	<p>UNIX: <code>plcclient.sh -s</code></p> <p>Windows: <code>plcclient -s</code></p>

Action	Command line
Dynamically change the log level of your data loader log files (temporarily).	UNIX: plcclient.sh -n loader_name -l log_level Windows: plcclient -n loader_name -l log_level where <ul style="list-style-type: none"> • <i>loader_name</i> is the name of your data loader • <i>log_level</i> is the log level of your data loader log files. If you restart the loader controller, these settings will revert back to the default level.
Dynamically change the log level of the log files for the reader or writer area of the specified data loader (temporarily).	UNIX: plcclient.sh -n loader_name -l log_level -a area Windows: plcclient -n loader_name -l log_level -a area where <ul style="list-style-type: none"> • <i>loader_name</i> is the name of your data loader • <i>area</i> is either reader or writer. • <i>log_level</i> is the log level of your data loader log files. If you restart the loader controller, these settings will revert back to the default level.

View or dynamically edit the data loader settings

Use the Platform Analytics Console to view or edit the data loader settings. Any changes you make to the settings are permanent (that is, even after restarting the loader controller).

1. In the navigation tree of the Platform Analytics Console, select Data Collection Nodes.
2. Right-click the loader controller for your cluster and select Loader properties.

Note:

You can only view the data loader properties when the corresponding loader controller is running.

3. Right-click the data loader you want to view or edit and select Properties.
4. Edit the data loader parameters, if needed.

You can edit the following data loader parameters:

- **Parameters:** The specific parameters for the data loader. You can only edit the parameters of FLEXnet data loaders (*flexnetcusageloader* and *flexnetceventsloader*).
 - **Interval (seconds):** The data gathering interval of the data loader, in seconds.
 - **Log level:** The data loader logs messages of a level specified here and higher.
 - **Reader Area:** The reader area of the data loader logs messages of a level specified here and higher. Specify *Inherit* to use the same log level as the entire data loader.
 - **Writer Area:** The writer area of the data loader logs messages of a level specified here and higher. Specify *Inherit* to use the same log level as the entire data loader.
 - **Description:** A description of the data loader.
5. To save any changes and close the window, click OK.

Data loaders

Platform Analytics node command-line tools

- [*dbconfig*](#) on page 42
- [*perfadmin*](#) on page 43
- [*plclient*](#) on page 44

dbconfig

Configures the node data source.

Synopsis

UNIX commands:

```
dbconfig.sh [add data_source_name | edit data_source_name]
```

```
dbconfig.sh -h
```

Windows commands:

```
dbconfig [add data_source_name | edit data_source_name]
```

```
dbconfig -h
```

Description

Run the command to configure the Platform Analytics node data source (ReportDB).

If you are running this command locally from an Platform Analytics node running UNIX, you need to be running X-Windows. If you are running this command remotely, you need to set your display environment.

Options

add *data_source_name*

Adds the specified data source to the Platform Analytics node.

edit *data_source_name*

Edits the specified data source on the Platform Analytics node.

-h

Prints the command usage and exits.

perfadmin

Administer the PERF services.

Synopsis

perfadmin start *service_name* | **all**

perfadmin stop *service_name* | **all**

perfadmin [**list** | **-h**]

Description

Starts or stops the PERF services, or shows status.

Run the command on the Platform Analytics node to control the loader controller service (pl c).

Options

start *service_name* | **all**

Starts the PERF services on the local host. You must specify the service name or the **all** keyword. Do not run this command on a host that is not the Platform Analytics node or the Platform Analytics server. You should only run one set of node services per cluster.

stop *service_name* | **all**

Stops the PERF services on the local host. You must specify the service name or the **all** keyword.

list

Lists status of PERF services. Run this command on the PERF host.

-h

Outputs command usage and exits.

Output

Status information and prompts are displayed in your command console.

SERVICE

The name of the PERF service.

STATUS

- **STARTED:** Service is running.
- **STOPPED:** Service is not running.
- **UNKNOWN:** Service status is unknown. The local host may not be the PERF host.

WSM_PID

Process ID of the running service.

HOST_NAME

Name of the host.

plcclient

Administer the loader controller or data loaders.

Synopsis

UNIX commands:

plcclient.sh [-s]

plcclient.sh [-l *log_level*]

plcclient.sh [-n *loader_name* -l *log_level*]

Windows commands:

plcclient [-s]

plcclient [-l *log_level*]

plcclient [-n *loader_name* -l *log_level*]

Description

Run the command to administer the loader controller or the data loaders.

Options

-s

View the status of the data loaders.

-l *log_level*

Dynamically change the log level of the loader controller to the specified log level. If you restart the loader controller (pl c) service, this setting will revert back to the default level.

-n *loader_name* -l *log_level*

Dynamically change the log level of the specified data loader to the specified log level. If you restart the loader controller (pl c) service, this setting will revert back to the default level.

Platform Analytics node configuration files

- [perf.conf](#) on page 46

perf.conf

The `perf.conf` file controls the operation of PERF.

About `perf.conf`

`perf.conf` specifies the version and configuration of various PERF components and features. The `perf.conf` file also specifies the file path to PERF directories and the PERF license file.

The `perf.conf` file is used by Platform Analytics and applications built on top of it. For example, information in `perf.conf` is used by Platform Analytics daemons and commands to locate other configuration files, executables, and services. `perf.conf` is updated, if necessary, when you upgrade to a new version of Platform Analytics.

Changing `perf.conf` configuration

After making any changes to `perf.conf`, run the following commands to restart the PERF services and apply your changes:

```
perfadmin stop all
```

```
perfadmin start all
```

Location

The default location of `perf.conf` is in `/conf`. If necessary, this default location can be overridden by modifying the `PERF_CONFDIR` environment variable.

Format

Each entry in `perf.conf` has the following form:

```
NAME=VALUE
```

The equal sign `=` must follow each `NAME` and there should be no space beside the equal sign. Text starting with a pound sign (`#`) are comments and are ignored. Do not use `#i f` as this is reserved syntax for time-based configuration.

DLP_ENABLED

Syntax

```
DLP_ENABLED=Y | N
```

Description

Enables data loss protection (DLP) for data loaders. If enabled, you can enable or disable data loss protection for specific data loaders in the Platform Analytics node by editing the specific data loader configuration file. If disabled, data loss protection is disabled in all data loaders in the Platform Analytics node and cannot be enabled in the specific data loader configuration file.

Default

Y (Enabled). In addition, all sampling data loaders have data loss protection enabled by default.

EGO_VERSION

Syntax

EGO_VERSION=*version_number*

Description

Specifies the version of EGO in the LSF cluster to which the Platform Analytics node belongs.

Example

```
EGO_VERSION=1.2
```

Default

By default, EGO_VERSION is set to the version of EGO in the LSF cluster to which the Platform Analytics node belongs.

LICENSE_FILE

Syntax

LICENSE_FILE="*file_name ... | port_number@host_name[:port_number@host_name ...]*"

Description

Specifies one or more demo or permanent license files used by Platform Analytics.

The value for LICENSE_FILE can be either of the following:

- The full path name to the license file.
 - UNIX example:


```
LICENSE_FILE=/usr/share/lsf/cluster1/conf/license.dat
```
 - Windows examples:


```
LICENSE_FILE=C:\licenses\license.dat
```

```
LICENSE_FILE=\\HostA\licenses\license.dat
```
- For a permanent license, the name of the license server host and TCP port number used by the lmgrd daemon, in the format *port@host_name*. For example:


```
LICENSE_FILE="1700@hostD"
```
- For a license with redundant servers, use a comma to separate the *port@host_names*. The port number must be the same as that specified in the SERVER line of the license file. For example:

UNIX:

```
LICENSE_FILE="port@hostA:port@hostB:port@hostC"
```

Windows:

```
LICENSE_FILE="port@hostA;port@hostB;port@hostC"
```

Multiple license files should be quoted and must be separated by a pipe character (|).

Windows example:

```
LICENSE_FILE="C:\licenses\license1|C:\licenses\license2|D:\mydir\license3"
```

Multiple files may be kept in the same directory, but each one must reference a different license server. When checking out a license, Platform Analytics searches the servers in the order in which they are listed, so it checks the second server when there are no more licenses available from the first server.

If this parameter is not defined, Platform Analytics assumes the default location.

Default

By default, `LICENSE_FILE` is set as the file path to the license file that you specified during the initial Platform Analytics installation.

If you installed FLEXlm separately from Platform Analytics to manage other software licenses, the default FLEXlm installation puts the license file in the following location:

- UNIX: `/usr/share/flexlm/licenses/license.dat`
- Windows: `C:\flexlm\license.dat`

LICENSE_VERSION

Syntax

`LICENSE_VERSION=version_number`

Description

Specifies the version of the license module installed with Platform Analytics.

Example

```
LICENSE_VERSION=7.0
```

Default

Not defined.

LOADER_BATCH_SIZE

Syntax

`LOADER_BATCH_SIZE=integer`

Description

Specifies the number of SQL statements that can be submitted to the database at the same time.

Valid values

Any positive, non-zero integer.

Default

5000

LSF_ENVDIR

Syntax

`LSF_ENVDIR=directory`

Description

Specifies the LSF configuration directory, which is the directory containing the `lsf.conf` file.

Default

/etc

LSF_VERSION

Syntax

LSF_VERSION=*version_number*

Description

Specifies the version of LSF in the cluster to which the Platform Analytics node belongs.

Example

```
LSF_VERSION=7.0
```

Default

By default, LSF_VERSION is set to the version of LSF in the cluster to which the Platform Analytics node belongs.

PERF_CONFDIR

Syntax

PERF_CONFDIR=*directory*

Description

Specifies the configuration directory, which contains the configuration files for Platform Analytics node components.

Default

- UNIX: *ANALYTICS_TOP*/conf
- Windows: *ANALYTICS_TOP*\conf

where *ANALYTICS_TOP* is the top-level Platform Analytics node installation directory.

PERF_LOGDIR

Syntax

PERF_LOGDIR=*directory*

Description

Specifies the logging directory, which contains the log files for Platform Analytics node components.

Default

- UNIX: *ANALYTICS_TOP*/log
- Windows: *ANALYTICS_TOP*\log

where *ANALYTICS_TOP* is the top-level Platform Analytics node installation directory.

PERF_TOP

Syntax

PERF_TOP=*directory*

Description

Specifies the top-level PERF directory.

Default

- UNIX: *ANALYTICS_TOP*
- Windows: *ANALYTICS_TOP*

where *ANALYTICS_TOP* is the top-level Platform Analytics node installation directory.

PERF_VERSION

Syntax

PERF_VERSION=*version_number*

Description

Specifies the version of PERF installed with the Platform Analytics node.

Example

```
PERF_VERSION=1. 2. 3
```

Default

Not defined.

PERF_WORKDIR

Syntax

PERF_WORKDIR=*directory*

Description

Specifies the working directory.

Default

- UNIX: *ANALYTICS_TOP/work*
- Windows: *ANALYTICS_TOP\work*

where *ANALYTICS_TOP* is the top-level Platform Analytics node installation directory

Managing Platform Analytics server

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

The server performs the following functions:

- Platform Analytics node management
- Cluster data management

12

Platform Analytics Console

The Platform Platform Analytics Console allows you to view cluster data and Platform Analytics configuration.

About the Platform Analytics Console

The Platform Platform Analytics Console displays information on your cluster and Platform Analytics configuration. You can also make some configuration changes to Platform Analytics components. You can view the following data in the Platform Analytics Console:

Clusters	Displays information on each cluster that Platform Analytics monitors.
Data Collection Nodes	This includes all Platform Analytics nodes in the system.
Data Sources	This includes the data sources that are running on the Platform Analytics server and nodes.
Scheduled Tasks	This includes the status and schedule of all scheduled tasks that the Platform Analytics server controls.
Events	Displays each event logged in Platform Analytics. You can filter the display of these events to find specific events.

Platform Analytics Console actions

Action	Command line
Start the Platform Analytics Console.	<p>UNIX: <i>ANALYTICS_TOP</i>/bin/runconsole.sh</p> <p>Windows: Start > Programs > Platform Platform Analytics Server > Platform Platform Analytics Console</p> <p>If you are running this command locally from the Platform Analytics server running UNIX, you need to be running X-Windows. If you are running this command remotely, you need to set your display environment.</p> <hr/> <p>Important:</p> <p>The Platform Analytics server must have access to the Platform Analytics data source (ReportDB). If the Platform Analytics server cannot connect to the data source, the data source configuration tool displays and the Platform Analytics Console will not start up until you can connect to the data source.</p>

13

Data transformers

Data transformers convert raw cluster data in the relational database into a format usable for reporting and analysis.

About data transformers

The LSF and license data is logged in the relational database in a raw format. At regular intervals, the data transformer converts this data to a usable format.

Logging levels

There are logging levels that determine the detail of messages that the data transformers record in the log files. In decreasing level of detail, these levels are ALL (all messages), TRACE, DEBUG, INFO, WARN, ERROR, FATAL, and OFF (no messages).

By default, the data transformers log messages of INFO level or higher (that is, all INFO, WARN, ERROR, and FATAL messages).

The data transformer log files are located in the `data transformer` subdirectory of your Platform Analytics server log directory:

- UNIX: `ANALYTICS_TOP/log/data transformer`
- Windows: `ANALYTICS_TOP\log\data transformer`

Default behavior


Data transformers convert data at a regular 10 minutes intervals. The following is a list of the data transformers and the database tables in which the data transformers generate the data:

Data transformer name	Transformed database tables
ClusterCapacity	RPT_CLUSTER_CAPACITY_RAW
FlexlmLicUsage	RPT_FLEXLM_LICUSAGE_RAW
Hardware	RPT_HARDWARE_RAW RPT_HARDWARE_DAY
WorkloadAccounting and Resource Usage	RPT_JOB MART_RAW RPT_JOB MART_DAY
WorkloadStatistics	RPT_WORKLOAD_STATISTICS_RAW

Data transformer interactions

Data transformers convert raw cluster data from the data tables through the server data sources in the relational database into a format usable for reporting and analysis.

The following diagram illustrates the interaction between the data transformers and other components.



Configuration to modify data transformer behavior

Action	Configuration files	Parameter and syntax
Specify the default log level of your data transformer log files.	<code>log4j.properties</code> File location: <code>ANALYTICS_TOP/conf</code> <code>log4j.properties</code>	<pre>log4j.appender.\${datatransformer} =log_level, \${datatransformer}</pre> <p>where</p> <ul style="list-style-type: none"> <code>log_level</code> is the default log level of your data transformer log files.
Specify the log level of the log file for the specified data transformer.		<pre>log4j.logger.transformer.datatransformer_name=log_level</pre> <p>where</p> <ul style="list-style-type: none"> <code>datatransformer_name</code> is the name of the data transformer. <code>log_level</code> is the log level of your data transformer log file. <p>For example, to set hardware to ERROR, add the following line to <code>log4j.properties</code>:</p> <pre>log4j.logger.transformer.hardware.loader=ERROR</pre>
Specify the log level of the log file for the Extractor or Loader in the ETL flow for the specified data transformer.		<pre>log4j.logger.transformer.datatransformer_name.component=log_level</pre> <p>where</p> <ul style="list-style-type: none"> <code>datatransformer_name</code> is the name of the data transformer. <code>component</code> is the ETL flow component. Use <code>extractor</code> to specify the Extractor and use <code>loader</code> to specify the Loader in the ETL flow. <code>log_level</code> is the log level of your data transformer Extractor or Loader log files. <p>For example, to set the Loader in WorkloadAccounting to WARN, add the following line to <code>log4j.properties</code>:</p> <pre>log4j.logger.transformer.WorkloadAccounting.loader=WARN</pre>

The data transformer only logs messages of the same or lower level of detail as `log_level`. Therefore, if you change the log level to ERROR, the data transformer will only log ERROR and FATAL messages.

Data transformer actions

Data transformers are installed as scheduled tasks. Change the schedule of data transformer services as you would for scheduled tasks (see [Scheduled tasks](#) on page 67).

14

Event notification

An event is a change in Platform Analytics reflecting a change in state.

About events

An event is a change in Platform Analytics reflecting a change in state, including events that provide information about problems encountered when running Platform Analytics (Warning, Error, or Fatal events), or events that contain useful administration information on Platform Analytics activities (Info events).

Event notification

Platform Analytics sends an event notification email when it encounters a change in state that matches the event notification settings. An event notification email informs you of the change in state in Platform Analytics or the cluster, allowing you to decide whether you want to check the Platform Analytics Console for further details.

Event actions

Action	Platform Analytics Console
View the list of events.	In the navigation tree, click Events .
View a filtered list of events.	When viewing the list of events, select Action > Filter Events from the menu toolbar.
Edit event notification settings.	When viewing the list of events, select Action > Notification from the menu toolbar.
<p>Important:</p> <p>If you enable or disable event notification, you need to restart the Platform Task Scheduler to apply this change. See Restart the Platform Task Scheduler on page 60.</p>	

Restart the Platform Task Scheduler

If you enable or disable event notification, you need to restart the Platform Task Scheduler to apply this change. The steps you take to restart the task scheduler depend on your operating system.

- Windows: Restart the task scheduler service.
 - a) From the Windows Control Panel, select Administrative Tools > Services.
 - b) Right-click Platform Platform Analytics Task Scheduler and select Restart.
- UNIX: Restart the task scheduler daemon.
 - a) From the command line, navigate to the `ANALYTICS_TOP/bin` directory.
 - b) Restart the Platform Analytics daemons.

perftool stop all

perftool start all

Configuration to modify event notification behavior

Action	Configuration files	Parameter and syntax
Filter specific event notification emails.	eventfilter.properties File location: <i>ANALYTICS_TOP/conf</i>	Add a new line for each filter. Email notifications that match any one of these lines are filtered out. Regular expressions are supported. For example, if the file contains the following: <pre> Communication timeout Connection reset PLC[0-9]+ has been restarted </pre> The following notifications will be filtered out and you will not receive these emails: Communication timeout PLC10 has been restarted at 12:00:00, Jan. 1, 2010.

Event notification

15

Data purger

The data purger (purger) service maintains the size of the database by purging old data from the database.

About the data purger

The relational database needs to be kept to a reasonable size to maintain optimal efficiency. The data purger manages the database size by purging old data from the database at regular intervals, which consists of dropping partitions that are older than the calculated data retention date.

Logging levels

There are logging levels that determine the detail of messages that the data loaders record in the log files. In decreasing level of detail, these levels are ALL (all messages), TRACE, DEBUG, INFO, WARN, ERROR, FATAL, and OFF (no messages).

By default, the data purger logs messages of ERROR level or higher (that is, all ERROR and FATAL messages) to the data purger log file, which is located in the Platform Analytics server log directory (*ANALYTICS_TOP/log* in the Platform Analytics server host).

Default behavior

The data purger runs as the following scheduled tasks on the Platform Analytics server:

- PartitionMaintenanceGroup1
- PartitionMaintenanceGroup2
- PartitionMaintenanceGroup3


Each scheduled task is responsible for purging different tables according to different schedules. This allows the workload to be split among different times.

Each scheduled task calculates the data retention date according to the data purger configuration, examines the tables (and their corresponding partitions) for which it is configured and drops any partitions that are older than the calculated data retention date.

Data purger interactions

The data purger drops database partitions from the data tables through the server data sources.

The following diagram illustrates the interaction between the data purger and other components.



Data purger actions

The data purger is installed as scheduled tasks. Change the schedules of the data purger services as you would for scheduled tasks (see [Scheduled tasks](#) on page 67).

Data purger

16

Scheduled tasks

Scheduled tasks are automated processing tasks that regularly run JavaScript-based scripts.

About scheduled tasks

After metric data is collected from hosts and stored in the database, the data undergoes several processing tasks for maintenance purposes. Platform Analytics automates the data processing by scheduling these processing tasks to run regularly. Each of these tasks calls a Javascript-based script.

You can modify these tasks, reschedule them, and create new scheduled tasks.

Scripts

Platform Analytics scheduled tasks call JavaScript-based scripts. These scripts work with data stored in the database for various maintenance tasks such as deleting old or duplicate records, or checking for problems with the collected data.

Predefined scheduled tasks

Platform Analytics includes several predefined scheduled tasks.

Data latency checker (`DataLatencyChecking`)

The data latency checker scheduled task checks the data latency in the data collected from the data loaders and data transformers. If the data latency is longer than the configured value or interval, the data latency checker sends an email notification.

By default, the data latency checker scheduled task runs every hour. If you want to modify the default configuration, edit `ANALYTICS_TOP/conf/health_check_notify.properties` and then restart the Platform Analytics server.

Daily Report (`DailyReportETL`)

Daily report builds jobmart data to `JOBMART_DAY` table and hardware data to `RPT_HARDWARE_DAY` table. By default, this task runs every day.

Cluster and Workload (`ClusterandWorkloadETL`)

Cluster and workload builds jobmart data to `RPT_CLUSTER_CAPACITY_RAW` table and hardware data to `RPT_WORKLOAD_STATISTICS_RAW` table. By default, this task runs every hour.

Hardware Jobmart (`HardwareJobmartLicusageETL`)

Hardware builds jobmart data to `RPT_JOBMART_RAW` table, hardware data to `RPT_HARDWARE_RAW` table and flexlm license usage data to `RPT_FLEXLM_LICUSAGE_RAW` table. By default, this task runs every hour.

(Vertica only) Data purger (`PartitionMaintenanceGroup*`)

The data purger scheduled tasks, which all have "PartitionMaintenanceGroup" in their names, control the data purger.

For more information, see [Data purger](#) on page 63.

(Vertica only) Duplicate record remover (`PKViolationClean`)

The duplicate record remover scheduler task checks the most recent data in the database (one to three days old) and deletes any duplicate records in the database (that is, those with a primary key violation).

This scheduled task is necessary because the Vertica database does not automatically delete records with a primary key violation.

By default, the duplicate record remover scheduled task runs every 12 hours.

(Oracle only) Tablespace partition (RawTablePartition and ReportTablePartition)

The tablespace partition scheduled tasks maintain partitioned tablespaces by creating new tablespaces or deleting tablespaces that are out of date. The RawTablePartition scheduled task maintains raw tablespace partitions and runs at 4:00 a.m. every day by default, while the ReportTablePartition scheduled task maintains aggregated tablespace partitions runs at 5:00 a.m. every day.

(Oracle only) Tablespace monitor (TSMonitor)

The TSMonitor scheduled task notifies you if any tablespaces have a utilization that exceeds a specified threshold. If you do not increase the tablespace size (or the tablespace does not increase in size automatically), the scheduled task sends a notification every hour. This scheduled task is disabled by default.

This scheduled task has the following prerequisites:

- To monitor the tablespaces, the Platform Analytics user on the Oracle database must have the appropriate privileges in the tablespaces (by selecting the privileges on the DBA_DATA_FILES and DBA_FREE_SPACE tables).
- Email notification must be enabled to receive the scheduled task notifications. See the Email notification chapter for more details.

Scheduled task actions

Action	Platform Analytics Console
View a list of scheduled tasks. You need to do this to perform any other action on the scheduled tasks.	In the navigation tree, click Scheduled Tasks .
Create a task in the list of scheduled tasks.	See Create, edit, or view a scheduled task on page 69 for detailed information.
View or edit a task from the list of scheduled tasks.	See Create, edit, or view a scheduled task on page 69 for detailed information.
Remove a task from the list of scheduled tasks.	In the main window, right-click the scheduled task and select Remove Scheduled Task .
Run a task manually from the list of scheduled tasks.	In the main window, right-click the scheduled task and select Run Now .

Create, edit, or view a scheduled task

Create, edit, or view a scheduled task.

You might edit a scheduled task for the following reasons:

- Schedule a task that is currently unscheduled
- Edit the next run time

- Edit the run interval
 - Add or edit task parameters
 - Modify how information about the task is logged and where it is stored
 - Modify the JavaScript file and function called by the task
1. In the navigation tree of the Platform Analytics Console, select Scheduled Tasks.
 2. Select the scheduled task to create, edit, or view.
 - To create a new scheduled task, right-click on the main window and select Add Scheduled Task.
 - To edit or view an existing scheduled task, right-click the scheduled task in the main window and select Edit Scheduled Task.

The Scheduled Task window for the scheduled task displays.

For an existing scheduled task, the following information is displayed in addition to the scheduled task parameters:

- Last Run Time: The previous time that this scheduled task was run.
 - Last Run Status: The status of the last run of this scheduled task.
 - Last Checkpoint: The last time the data was checkpointed during the scheduled task. If the checkpoint and the scheduled task are completed, this is "DONE".
3. Edit the scheduled task parameters that you want to change.

Caution:

Do not change the name of the scheduled task; otherwise, Platform Analytics may have problems with scheduling your renamed task.

- a) To change the script file for the task, specify the new script file in the Script File field.

The script file must reside in the *ANALYTICS_TOP* directory. If it is in a subdirectory, include the file path of the subdirectory in the field.

For example, if the new script file is *new_script.js* and resides in the *ANALYTICS_TOP/bin* directory, define the new script file as the following:

```
/bin/new_script.js
```
- b) To change the function to run in the script for the task, specify the new script function in the Script Function field.

The script can include other functions, but the other functions will run only if they are called by this specified script function.
- c) To change the log file for this task, specify the new log file in the Log File field.

The location of the log directory is as follows:

 - UNIX: *ANALYTICS_TOP/log*
 - Windows: *ANALYTICS_TOP\log*
- d) To change the level of detail of information recorded in the log file, select the new log level in the Log Level field.

All messages of this level or lower are recorded in the log file. In decreasing level of detail, the logging levels are DEBUG, VERBOSE, INFO, WARNING, and ERROR.

For example, if you specify "INFO", the log file contains INFO, WARNING, and ERROR messages.
- e) To enable scheduling for this task, enable the Enable Scheduling check box.

- f) To change the next date and time that this task is scheduled to run, modify the fields in the Next Run Time box.
- g) To change the run interval of the scheduled task to a fixed interval, select the Run every: field and specify the interval.
- h) To change the run interval of the scheduled task to a calculated value, select the Call this function field specify the function in the script file to determine the run interval.

The function must return a time stamp string in the following format:

```
YYYY-MM-DD hh: mm: ss. xxxx
```

This time stamp indicates the the next date and time in which this task is scheduled to run.

- i) To add optional parameters that Platform Analytics looks for in the script file, enter them into the Parameters field.

This field does not exist in certain scheduled tasks.

4. For the TSMonit or scheduled task, specify the details of the tablespaces you want to monitor.

- a) Specify the name of the data source in the Data Source field.

The default name of the data source is **ReportDB**.

- b) Specify a comma-separated list of the tablespaces for TSMonit or to monitor in the Tablespaces field. Leave this field blank if you want TSMonit or to monitor all tablespaces in the Oracle database.

For example, to monitor the TS_DATA_01, TS_DATA_02, and SYSTEM tablespaces,

TS_DATA_01, TS_DATA_02, SYSTEM

- c) Specify a threshold for tablespace utilization, as a percentage or a decimal value between 0 and 1, in the Threshold field.

If there is at least one monitored tablespace that exceeds this threshold, TSMonit or sends a notification every hour until the tablespaces no longer exceed the threshold.

For example, to set a threshold of 90%, specify **0.9** or **90%** in the Threshold field.

5. To save your changes and close the window, click OK.

Scheduled tasks

Platform Analytics server command-line tools

- [*perfadmin*](#) on page 74
- [*runconsole*](#) on page 76

perfadmin

Administer the PERF services.

Synopsis

perfadmin start *service_name* | **all**

perfadmin stop *service_name* | **all**

perfadmin [**list** | **-h**]

Description

Starts or stops the PERF services, or shows status.

Run the command on the Platform Analytics server to control the task scheduler service (*pat s*) and the remoting server service (*par s*, if the asynchronous data loading mode is enabled).

Options

start *service_name* | **all**

Starts the PERF services on the local host. You must specify the service name or the *all* keyword. Do not run this command on a host that is not the Platform Analytics node or the Platform Analytics server. You should only run one set of node services per cluster.

stop *service_name* | **all**

Stops the PERF services on the local host. You must specify the service name or the *all* keyword.

list

Lists status of PERF services. Run this command on the PERF host.

-h

Outputs command usage and exits.

Output

Status information and prompts are displayed in your command console.

SERVICE

The name of the PERF service.

STATUS

- **STARTED:** Service is running.
- **STOPPED:** Service is not running.
- **UNKNOWN:** Service status is unknown. The local host may not be the PERF host.

WSM_PID

Process ID of the running service.

HOST_NAME

Name of the host.

runconsole

Starts the Platform Analytics console.

Synopsis

runconsole.sh

runconsole

`runconsole.sh` is the command for UNIX and `runconsole` is the command for Windows.

If you are running this command locally from the Platform Analytics server running UNIX, you need to be running X-Windows. If you are running this command remotely, you need to set your display environment.

Platform Analytics server configuration files

- [pi.conf](#) on page 78

pi.conf

The `pi.conf` file controls the operation of the Platform Analytics server.

About pi.conf

`pi.conf` specifies the configuration of various Platform Analytics server components and features.

Changing pi.conf configuration

After making any changes to `pi.conf`, run the following commands from the `ANALYTICS_TOP/bin` directory to restart the Platform Analytics server and apply your changes:

```
perfadmin stop all
```

```
perfadmin start all
```

Location

The location of `pi.conf` is in `ANALYTICS_TOP/conf`.

Format

Each entry in `pi.conf` has the following form:

```
NAME=VALUE
```

The equal sign `=` must follow each `NAME` and there should be no space beside the equal sign. Text starting with a pound sign (`#`) are comments and are ignored. Do not use `#if` as this is reserved syntax for time-based configuration.

PIAM_PORT

Syntax

```
PIAM_PORT=port_number
```

Description

Specifies the Platform Automation Manager listening port number.

Default

9991

CHECK_INTERVAL

Syntax

```
CHECK_INTERVAL=time_in_seconds
```

Description

Specifies the interval, in seconds, that the Platform Automation Manager checks the system.

Default

60 seconds

send_notifications

Syntax

send_notifications=true | false

Description

Enables event notification.

You would normally configure this parameter using the Platform Analytics Console (in the navigation tree, click Events, then right-click on the list of events and select Action > Notification).

If set to true, Platform Analytics sends an event notification email when it encounters a change in state that matches the event notification settings. An event notification email informs the you of the change in state in Platform Analytics or the cluster, allowing you to decide whether you want to check the Platform Analytics Console for further details.

For more information on event notification, refer to [Event notification](#) on page 59.

Default

true

mail.smtp.host

Syntax

mail.smtp.host=host_name.domain_name

Description

Specifies the SMTP server that Platform Analytics uses to send event notification emails.

You would normally configure this parameter using the Platform Analytics Console (in the navigation tree, click Events, then right-click on the list of events and select Action > Notification).

Example

```
mail.smtp.host=smtp.example.com
```

Valid values

Any fully-qualified SMTP server name.

Default

Not defined.

from_address

Syntax

from_address=email_account

Description

Specifies the sender email address that Platform Analytics uses to send event notification emails.

You would normally configure this parameter using the Platform Analytics Console (in the navigation tree, click Events, then right-click on the list of events and select Action > Notification).

Example

```
from_address=system@example.com
```

Default

Not defined

to_address

Syntax

```
to_address=email_account
```

Description

Specifies the email addresses of the intended recipient of the event notification emails that Platform Analytics will send.

You would normally configure this parameter using the Platform Analytics Console (in the navigation tree, click Events, then right-click on the list of events and select Action > Notification).

Example

```
to_address=admin@example.com
```

Default

Not defined

subject_text

Syntax

```
subject_text=text
```

Description

Specifies the subject of the event notification emails that Platform Analytics will send.

You would normally configure this parameter using the Platform Analytics Console (in the navigation tree, click Events, then right-click on the list of events and select Action > Notification).

Example

```
subject_text=Platform Platform Analytics Error Notification
```

Default

Not defined

message_header

Syntax

```
message_header=text
```

Description

Specifies the header of the event notification emails that Platform Analytics will send. The rest of the email contains information about the event change and is not specified here.

You would normally configure this parameter using the Platform Analytics Console (in the navigation tree, click Events, then right-click on the list of events and select Action > Notification).

Example

```
message_header=An error has occurred in the Platform Platform Analytics data collection system.
```

Default

Not defined

PIEM_PORT

Syntax

PIEM_PORT=*port_number*

Description

Specifies the Platform Event Manager listening port number.

Default

37600

PIEM_HOST

Syntax

PIEM_HOST=*port_number*

Description

Specifies the Platform Event Manager host.

Default

local host

PIEM_TIMEOUT

Syntax

PIEM_TIMEOUT=*time_in_seconds*

Description

Specifies the timeout, in seconds, for Platform Event Manager to receive events.

Default

36000 seconds (10 hours)

EVENTLOGGER_TIMEOUT

Syntax

EVENTLOGGER_TIMEOUT=*time_in_seconds*

Description

Specifies the timeout, in seconds, for the Platform Event Manager client to send event notifications.

Default

5 seconds

EVENT_LEVEL

Syntax

EVENT_LEVEL=**ALL** | **TRACE** | **DEBUG** | **INFO** | **WARN** | **ERROR** | **FATAL** | **OFF**

Description

Specifies the logging levels of events to send to the Platform Event Manager. All events of this specified level or higher are sent. In decreasing level of detail, these are TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.

Use ALL to specify all messages and OFF to specify no messages.

Example

```
EVENT_LEVEL=WARN
```

All WARN, ERROR, and FATAL messages are sent to Platform Event Manager.

Default

INFO

All INFO, WARN, ERROR, and FATAL messages are sent to Platform Event Manager.

DS_NAME

Syntax

DS_NAME=*data_source_name*

Description

Specifies the name of the data source for the Platform Event Manager to access.

Default

ReportDB

PURGER_BATCH_SIZE

Syntax

PURGER_BATCH_SIZE=*integer*

Description

Specifies the number of records to purge in each batch.

Valid values

Any positive integer

Default

10000000

SHOW_BUSINESS_INFO

Syntax

SHOW_BUSINESS_INFO=YES | Y | NO | N

Description

Specify YES or Y to enable the Data Collection Nodes page in the Platform Analytics Console to display the following optional columns:

- System Purpose
- Display Description
- Business Area

Default

YES

Platform Analytics server configuration files



Viewing reports

The support hosts such as Analytics reporting server, Analytics Designer, and Platform Application Center do not run Platform Analytics, they are necessary in order for you to take full advantage of the cluster operations data and reports that Platform Analytics assembles and generates.

Viewing reports

Generating reports

The Platform Analytics reporting server generates Platform Analytics reports and allows other users to view these reports.

The Platform Analytics reporting server runs Tableau Server, which is a ROLAP (Relational Online Analytics Processing) analytic tool for business intelligence that provides browser-based reports. The Platform Analytics reporting server uses Tableau Server to generate the Platform Analytics reports and allows other users to view these reports.

The Platform Analytics reporting server may run on the same host as the Platform Analytics server if that host meets the Tableau Server system requirements.

The Platform Analytics reporting server provides the following default workbooks to allow you to analyze your clusters:


Workbook name	Description
Workload Statistics	Reports information about all jobs in any state that are sampled from all active LSF clusters. This allows you to perform a detailed analysis of current LSF workload at any time period.
Workload Accounting	Reports job information from LSF job finish events. This allows you to perform a detailed analysis of completed LSF jobs in all clusters.
License Usage	Reports FlexNet Server license usage on any license server or across multiple license servers. This allows you to analyze the usage, consumption, and utilization of licenses by users and hosts.
Hardware	Reports the hardware utilization at any time period.
Cluster Capacity	Reports the usage of all slots in LSF and the workload being run. This allows you to identify IDLE, DOWN, CLOSED, and RUNNING capacity.
Workload Accounting (Daily) and Hardware (Daily)	Data is aggregated daily for better workbook performance.
Resource Memory Requested Vs Used	Reports wasted memory usage information by comparing required and used memories.

If you want to modify or create a new report, use the Platform Analytics designer.

Platform Analytics reporting server interactions

The Platform Analytics reporting server obtains time series data from the database through the Tableau Server data sources. All data obtained by the Platform Analytics reporting server are assembled into reports are then accessible from the Platform Application Center.

The following diagram illustrates the interaction between the Platform Analytics reporting server and other components.



Collecting data and viewing reports

The Platform Analytics reporting server generates Platform Analytics reports and allows other users to view these reports. In order to view reports, you need to first collect data, publish them to Analytics reporting server, and view them.

Collecting data

If you want to collect FLEXlm usage and FLEXlm events data, start the license servers and configure the Analytics node.

1. Start the LSF cluster.

Run `lsfstartup` after sourcing the `lsf.profile` file.

2. Start the license server daemon.

a) Log on to the license server host as LSF administrator.

b) Run the `lmgrd` command in `LSF_SERVERDIR` to start the license server daemon:

```
lmgrd -c /usr/share/lsf/lsf_62/conf/license.dat -l /usr/share/lsf/lsf_62/log/license.log
```

c) Make sure that the FLEXnet data loaders are enabled in your cluster.

3. Start the database.

1. Open the Administration Tools.

2. On the Main Menu, select Start Database.

4. Start the Platform Analytics node and source LSF and perf environment.

perfadmin start plc | all

plcclient [-s]

Check the `plc` configuration file for any errors `plc.log.<host_name>` under the `ANALYTICS_TOP/log` directory.

Check log file of individual loaders (`<data_loader_name>.log.<host_name>`) under the `ANALYTICS_TOP/log/data_loader` directory for details of individual data loaders.

You can even check the database table to see if data has been successfully loaded into the database.

5. Start the Platform Analytics server and transform data

perfadmin start all

runconsole

Check log files under the `ANALYTICS_TOP/log` directory for details.

Viewing reports

Once data is collected in the database, you can view reports using the Analytics reporting server. Optionally, you can even view reports using Platform Analytics Designer or Platform Application Center.

1. Log in to the Platform Analytics reporting server.

`http://<host_name>:<port>`

where *<host_name>* is the name of the system where Tableau server is installed and *<port>* is the number which you entered during the Tableau server installation.

2. You can view workbooks, worksheets, and dashboards.
 - Workbook—A Workbook is a Tableau report (twb) file. It consists of Dashboards and Worksheets.
 - Dashboard—A Dashboard is a view of multiple Worksheets.
 - Worksheet—A Worksheet is a single view of queried data from a data source. This may be a table or a chart. A worksheet does not have to be viewed via a dashboard, it can be accessed directly if required.

Collecting data and viewing reports

Platform Application Center (optional)

Platform Application Center allows users and administrators to monitor hosts and to submit and monitor jobs.

About the Platform Application Center host


The Platform Application Center host communicates with the Platform Analytics reporting server to provide access to the Platform Analytics reports to monitor the Platform LSF clusters. It also provides browser-based access to all compatible Platform applications.

For more details, refer to the Platform Application Center documentation.

Platform Application Center host interactions

The Platform Analytics reporting server obtains time series data from the database through the Tableau Server data sources. All data that the Platform Analytics reporting server obtains and assembles into reports are then accessible from the Platform Application Center.

The following diagram illustrates the interaction between the support hosts and other components.



Enable HTTPS for Platform Application Center and Tableau

Follow the steps below to configure HTTPS on both Platform Application Center and Tableau using a self-signed certificate.

These instructions apply to:

- Platform Application Center 8.0.1, 8.0.2
- Platform Analytics 8.0, 8.0.2, with Tableau 6.0, 6.1

Note that you can configure HTTPS only for Platform Application Center, only for Tableau, or for both.

When you configure HTTPS for Platform Application Center, it affects access to the web server (URL will be https), access to Web Services, and the Report Builder (Report Builder will need a certificate to communicate with Platform Application Center).

When you configure HTTPS for Tableau, it affects report generation and Workbook access.

Enable HTTPS for Platform Application Center

The following steps use Platform Application Center's own self-generated certificate to enable HTTPS. If you want to use your own certificate, the same steps may not apply.

1. Log in to the Platform Application Center web server as root.
2. Set your Platform Application Center environment.

For example:

- For csh or tcsh:

```
# source /opt/pac/cshrc.platform
```
- For sh, ksh, or bash:

```
# ./opt/pac/profile.platform
```

3. Enable HTTPS and restart Platform Application Center.

```
# pmcadmin https enable  
# pmcadmin stop  
# pmcadmin start
```

4. Generate the certificate file.

For example:

```
/opt/pac/jre/linux-x86_64/bin/keytool -export -alias tomcat -file server.crt
```

Note:

You will be prompted for a password. Enter "changeit".

This command generates a file named `server.crt`. Copy this file to a temporary directory on the Analytics reporting server.

5. Log in to the Analytics reporting server as the Tableau administrator.
6. Import the `server.crt` file that you generated in step 4.

For example:

```
C:\analytics8.0_reports\jre\bin\keytool -import -file server.crt -keystore C:\analytics8.0_reports\jre\lib\security\cacerts
```

Note:

You will be prompted for a password. Enter "changeit".

- Edit the Platform Analytics Report Builder configuration file `c:\analytics8.0_reports\conf\rptbuilder.conf` and change the `PACServerUrl` to HTTPS.

For example:

```
PACServerUrl=https://192.168.0.1:8443
```

- Restart the Platform Analytics Report Builder service.

For example:

```
C:\analytics8.0_reports\bin\perfadmin.bat stop parb  
C:\analytics8.0_reports\bin\perfadmin.bat start parb
```

- Test that HTTPS is working by trying to access Platform Application Center with a web browser using `https://`.

Enable HTTPS for Tableau

- Log in to the Tableau server as local administrator.
- Create an SSL certificate and key for Tableau:

The following step is reproduced from the Tableau Knowledge Base (<http://www.tableausoftware.com/support/knowledge-base/creating-ssl-certificate-and-key-tableau-server>)

- To create a key

Step 1

Open the Command Prompt, and change directories to the path specified below, based on your operating system:

- On a 32-bit machine: `C:\Program Files\Tableau\Tableau Server\version\apache\bin`
- On a 64-bit machine: `C:\Program Files (x86)\Tableau\Tableau Server\version\apache\bin`

Step 2

Execute the command `openssl.exe genrsa -des3 -out yourcertname.key 4096` from the Command Prompt to create your key file.

Note: This command uses a 4096 bit modulus for the key. Other values, such as 1024 bits can be used, but provides less security. If a value is not provided, 512 bits is used.

Step 3

Type a passphrase after being prompted.

Important: Do not forget this passphrase.

Step 4

Execute the command `openssl.exe rsa -in yourcertname.key -out yourcertname.key` from the Command Prompt to embed your passphrase.

Note: Although embedding a passphrase may compromise the security of the certificate, Tableau Server requires that the passphrase is embedded.

- To generate a CSR

Step 1

Execute the command `openssl.exe req -new -key yourcertname.key -out yourcertname.csr` command from the Command Prompt to create the CSR file.

Note: If you see an error message about the config information being unable to load, retype the command above with `-config ..\conf\openssl.cnf`. Alternatively, you can set an environment variable to resolve the issue by typing the following command:

```
set OPENSSL_CONF=c:\Program Files\Tableau\Tableau Server\6.0\apache\conf\openssl.cnf
```

Step 2

Enter the required information after being prompted.

Note: When prompted to enter the Common Name value, type in the server name. If the common name and server name are different, errors will occur when a browser or Tableau Desktop try to connect to the server.

Step 3

Create a certificate by sending it to a commercial provider or by signing it yourself.

Once you have the key and certificate file, you can apply it to Tableau Server using the instructions in the "Configuring SSL" section of the Tableau Server Administrator Guide (<http://www.tableausoftware.com/currentadmin.php>).

3. Create a self-signed certificate.

```
openssl.exe x509 -req -days 365 -in yourcertname.csr -signkey yourcertname.key -out tableau.crt
```

4. Apply the private key and certificate to Tableau

- a) Log in to the Analytics Reporting Server as the user under which the Tableau Server service is running.
- b) Shut down the Tableau Server service from the Windows Services Controller.
- c) Select Start > Platform Analytics Server 6.0 > Configure Platform Analytics Server.

The Tableau Server Configuration dialog box is displayed.

- d) Select the SSL tab and configure SSL settings.

The following information is reproduced from the *Tableau Administrator Guide, Configuring SSL* chapter (http://downloads.tableausoftware.com/quickstart/server-guides/server_admin6.0.pdf)

Select the option to Use SSL for Server Communication. Then specify a location for each of the following certificate files. These files should be located on the local machine.

- SSL Certificate File - must be a valid PEM encoded x509 certificate with the extension `.crt`
- SSL Certificate Key File - must be a valid RSA or DSA key that is not password protected with the file extension `.key`
- SSL Certificate Chain File (Optional) - Some certificate providers issue two certificates for Apache. The second certificate is the chain file that contains information about the provider. If your provider has issued this second certificate you can enter it here.

When finished, click OK.

The changes will take effect the next time the server is restarted. When the server is configured for SSL, it will accept requests to the non-SSL port (default is port 80) and automatically redirects to the SSL port 443.

SSL errors are logged in the install directory at the following location. Use this log to troubleshoot validation and encryption issues.

C:\ProgramData\Tableau\Tableau Server\data\tabsvc\logs\httpd\error.log

NOTE: Tableau Server only supports port 443 as the secure port. It cannot run on a machine where any other application is using port 443.

5. Start the Tableau Server service from the Windows Services Controller.
6. Configure the Report Builder.
 - a) Log in to the Analytics Reporting Server as the local administrator.
 - b) Import the public key(certificate) for Report Builder , replacing REPORT_JRE_HOME with your own path.

Note:

You will need to provide an alias to avoid conflict with the default alias "mykey" for the Platform Application Center certificate.

```
C:\analytics8.0_reports\jre\bin\keytool -import -file tableau.crt -alias tableau -keystore REPORT_JRE_HOME/lib/security/cacerts
```

- c) Enter the keystore password.
- d) Edit analytics8.0_reports\conf\rptbuilder.conf and set these parameters to the following values:

```
TableauSSLEnabled=Y
AnalyticsReportingServerPort=443
```

- e) Restart the Platform Analytics Report Builder service.

For example:

```
C:\analytics8.0_reports\bin\perfadmin.bat stop parb
C:\analytics8.0_reports\bin\perfadmin.bat start parb
```

7. Configure Platform Application Center.
 - a) Log in to the Platform Application Center web server as root.
 - b) Import the public key(certificate) for Platform Application Center, replacing /opt/pac with the directory in which you installed Platform Application Center.

```
# keytool -import -file tableau.crt -alias tableau -keystore /opt/pac/jre/linux-x86_64/bin/keytools
```

- c) Set your Platform Application Center environment:

For example:

- For csh or tcsh:


```
# source /opt/pac/cshrc.platform
```
- For sh, ksh, or bash:


```
# . /opt/pac/profile.platform
```

- d) Edit /opt/pac/gui/conf/pmc.conf and specify the Tableau server host name.

```
TABLEAU_SERVER=https://Tableau_host_name:443
```

For example:

```
TABLEAU_SERVER=https://tabv6.lsf.platform.com:443
```

Platform Application Center (optional)

- e) Restart Platform Application Center to apply the changes.

```
# pmcadmin stop  
# pmcadmin start
```

- 8. Install the SSL certificate (.crt file) on the Tableau Desktop.
 - a) Log in to Windows as a domain or local administrator.
 - b) Select Start > Run, type mmc and press Enter.
 - c) Select File > Add/Remove Snap-ins.
 - d) Under Available snap-in select Certificates and click Add.
 - e) Select Computer account and click Next.
 - f) Leave Local computer checked and click Finish.
 - g) Expand Certificates (Local Computer) > Trusted Root Certification Authorities > Certificates.
 - h) Right-click Certificates and select All Tasks > Import.
 - i) Browse to the .crt certificate file used by the Tableau Server, and select Open.
 - j) Click Next.
 - k) Ensure that the Certificate Store field has the value Trusted Root Certification Authorities, and click Next.
 - l) Click Finish, a popup message is displayed, click OK.

P A R T

VI

Managing Platform Analytics

Secure your data and working environment

Customize the security of your cluster to secure your data and working environment.

Actions to secure your data and working environment

- [Open ports to communicate across firewalls](#) on page 104
- [Modify the database password](#) on page 104

Open ports to communicate across firewalls

If your cluster extends across the Internet securely, the server has to communicate with other hosts in the cluster across firewalls. Platform Analytics uses the following ports to communicate with other hosts in the cluster:

Port name	Default port number	Additional information
PIEM_PORT	9091	Internal port for the event manager. Used for receiving events from Platform Analytics components. Configuration is not required,
PIAM_PORT	9092	Internal port for the automation manager. Used for receiving events from Platform Analytics components. Configuration is not required.
Remoting server port (asynchronous data loading mode only)	9093	Internal port for the remoting server. Used for communicating between the remoting server and the remoting node. Configuration is not required. This port is only used if you enabled the asynchronous data loading mode.

1. Edit *ANALYTICS_TOP/conf/pi.conf* to open the appropriate ports.
2. Restart the Platform Analytics Console to start communicating with the new ports.

Modify the database password

If you modify the password that Platform Analytics data sources use to connect to the database, you must update Platform Analytics to use the new password.

1. Log into the Platform Analytics Console.
2. In the navigation tree, select Data Sources.
3. In the right pane, right-click ReportDB and select Edit Data Source.

The Data Source Properties window displays.

4. Specify the new password.
5. To verify the database connection, click Test.
6. To save your changes, click OK.

Maintaining the database

This section describes the relevant parts in the *Administrator's Guide* for the Vertica Analytic Database that you need to refer to for more details on maintaining the database. All of the following sections are located in the *Operating the Database* chapter of the *Administrator's Guide* for the Vertica Analytic Database.

Actions to maintain the database

- Partition tables in the database.

You can partition data tables in the Vertica database, which divides one large table into smaller tables. This can optimize query performance by utilizing parallel performance of the disks in which the table partitions reside.

For more details on recovering the database, refer to *Partitioning Tables* in the *Administrator's Guide* for the Vertica Analytic Database.

- Recover the database.

You can recover the database to a functional state after at least one node in the system fails.

For more details on recovering the database, refer to *Recovering the Database* in the *Administrator's Guide* for the Vertica Analytic Database.

- Back up or restore data in the database.

You can back up or restore data in the database using full backups or incremental backups. You can use backups to recover a previous version

Back up and restore data in the database

You can back up or restore data in the database using full backup or incremental backup scripts.

Back up or restore data in the database before you

- upgrade Vertica to a newer version,
- drop a partition,
- add a node to your database cluster.

Attention:

Following are some of the important points that you have to remember before proceeding to back up your files:

- Make sure you have installed `rsync 3.0` or later on the database nodes. You can use `rsync --version` to check the version.
 - Check the disk space in every vertica nodes and make sure that the backup directory has enough space.
 - The `backup.sh` script works only if the database is up and running. You can use `admin` tools in Vertica to check the database status.
 - It is important to note the snapshot name used by the `backup.sh` script for use in restore operations.
 - It is recommended to do a full backup at least once a week and incremental backup every other day.
-

Full backup

You can either use cold backup or hot backup to back up all the data on the drive.

- Cold backup—this is an offline backup. Make sure that the database is down before you copy all data to a backup directory.
- Hot backup—this is a dynamic backup. Vertica provides a utility to perform full backup called `backup.sh`.

For more information on backing up or restoring data in the database, refer to Backup and Restore in the *Administrator's Guide* for the Vertica Analytic Database

Incremental Backup

You can do an incremental backup to back up files that have changed or are new since the last incremental backup. This method takes less time to back up data compared to full backup.

1. Do a hot backup first. Vertica creates a snapshot file. This file is found in the location where you set `-B` parameter while you use `backup.sh` to full backup your database.
2. Do the incremental backup, you can also use `backup.sh` in Vertica(`$vertica_top/scripts/`). You must specify the snapshot file that was created by full backup. Refer to Backup and Restore in the *Administrator's Guide* for the Vertica Analytic Database for more details.

You can write a script to run incremental backup every other day. For example, `/opt/Vertica/scripts/backup.sh -s host1,host2,host3 -i host1 -b host1 -B /backupDir -D /vdata/pa8 -d pa8 -u dbadmin -w dbadmin -S backup1`

This creates a backup from a 3 nodes system and is run from `host1` and initialized by `host1` and the backup is stored under `/backupDir`.

Restore

Attention:

Following are some of the important points to note before restoring your files:

- The backup must have been created using the `backup.sh` script. It is important to note the snapshot name used by the `backup.sh` script for use in restore operations.
- By default, `restore.sh` does not restore the `vertica.conf` file. This is useful if you have modified the database configuration since the database was backed up. Use the `restore.sh` script with the `-c` switch to restore the `vertica.conf` file.

For example, `restore.sh -c`

- Make sure that you shutdown the database before running the restore script.
-

Use `restore.sh (/opt/vertica/scripts)` script to restore the database from the backup created by `backup.sh`.

For example, `/opt/Vertica/bin/restore.sh -s host1, host2, host3 -b host1 -B /backupDir -D /vdata/pa8 -S backup1`

This restores snapshot `backup1` to a 3 node system from backup directory `/backupDir` from backup host, `host1`.

24

Troubleshooting the node

Actions to troubleshoot the Platform Analytics node

- [Change the default log level of your log files](#) on page 110
- [Disable data collection for individual data loaders](#) on page 110
- [Check the status of the loader controller](#) on page 111
- [Check the status of the data loaders](#) on page 112
- [Check the status of the Platform Analytics node database connection](#) on page 112
- [Check core dump on the Platform Analytics node](#) on page 112
- [Debug LSF API](#) on page 115
- [Analytics node did not respond](#) on page 115

Change the default log level of your log files

Change the default log level of your log files if they do not cover enough detail, or cover too much, to suit your needs.

1. If you are 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**
2. If you are logged into a UNIX host, source the PERF environment.
 - For `csh` or `tcsh`: **source \$PERF_TOP/conf/cshrc.perf**
 - For `sh`, `ksh`, or `bash`: **.\$PERF_TOP/conf/profile.perf**

3. Edit the `log4j.properties` file.

This file is located in the PERF configuration directory:

- UNIX: `$PERF_CONFDIR`
- Windows: `%PERF_CONFDIR%`

4. Navigate to the section representing the service you want to change, or to the default loader configuration if you want to change the log level of the data loaders, and look for the `*.logger.*` variable.

For example, to change the log level of the loader controller log files, navigate to the following section, which is set to the default `INFO` level:

```
# Loader controller ("plc") configuration
log4j.logger.com.platform.perf.dataloader=INFO com.platform.perf.dataloader
```

5. Change the `*.logger.*` variable to the new logging level.

In decreasing level of detail, the valid values are `ALL` (for all messages), `DEBUG`, `INFO`, `WARN`, `ERROR`, `FATAL`, and `OFF` (for no messages). The services or data loaders only log messages of the same or lower level of detail as specified by the `.logger.*` variable. Therefore, if you change the log level to `ERROR`, the service or data loaders will only log `ERROR` and `FATAL` messages.

For example, to change the loader controller log files to the `ERROR` log level:

```
# Loader controller ("plc") configuration
log4j.logger.com.platform.perf.dataloader=ERROR com.platform.perf.dataloader
```

6. Restart the service that you changed (or the loader controller if you changed the data loader log level).

Disable data collection for individual data loaders

To reduce unwanted data from being logged in the database, disable data collection for individual data loaders.

1. If you are logged into a UNIX host, source the LSF environment.
 - For `csch` or `tcsh`: **source `LSF_TOP/conf/cshrc.lsf`**
 - For `sh`, `ksh`, or `bash`: **. `LSF_TOP/conf/profile.lsf`**
2. If you are logged into a UNIX host, source the PERF environment.
 - For `csch` or `tcsh`: **source `$PERF_TOP/conf/cshrc.perf`**
 - For `sh`, `ksh`, or `bash`: **. `$PERF_TOP/conf/profile.perf`**
3. Edit the `plc` configuration files for your data loaders.
 - For host-related data loaders, edit `plc_ego.xml` and `plc_coreutil.xml`.
 - For job-related data loaders (LSF data loaders), edit `plc_lsf.xml` and `plc_bj_obs-sp012.xml`.
 - For advanced job-related data loaders (advanced LSF data loaders), edit `plc_lsf_advanced_data.xml`.
 - For license-related data loaders (FLEXnet data loaders), edit `plc_license.xml`.

These files are located in the LSF environment directory:

- UNIX: `$LSF_ENVDIR`
- Windows: `%LSF_ENVDIR%`

4. Navigate to the specific `<DataLoader>` tag with the `Name` attribute matching the data loader that you want to disable.

For example:

```
<DataLoader Name="hostgrouploader" ... Enable="true" ... />
```

5. Edit the `Enable` attribute to `"false"`.

For example, to disable data collection for this plug-in:

```
<DataLoader Name="hostgrouploader" ... Enable="false" ... />
```

6. Restart the `plc` service.

Check the status of the loader controller

1. If you are logged into a UNIX host, source the LSF environment.
 - For `csch` or `tcsh`: **source `LSF_TOP/conf/cshrc.lsf`**
 - For `sh`, `ksh`, or `bash`: **. `LSF_TOP/conf/profile.lsf`**
2. If you are logged into a UNIX host, source the PERF environment.
 - For `csch` or `tcsh`: **source `$PERF_TOP/conf/cshrc.perf`**
 - For `sh`, `ksh`, or `bash`: **. `$PERF_TOP/conf/profile.perf`**
3. Navigate to the PERF binary directory.
 - UNIX: **cd `$PERF_TOP/version_number/bin`**
 - Windows: **cd `%PERF_TOP%\version_number\bin`**
4. View the status of the loader controller (`plc`) and other PERF services.

perfadmin list
5. Verify that there are no errors in the loader controller log file.

The loader controller log file is located in the log directory:

- UNIX: `$PERF_LOGDIR`

- Windows: %PERF_LOGDIR%

Check the status of the data loaders

1. If you are 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**
2. If you are logged into a UNIX host, source the PERF environment.
 - For csh or tcsh: **source \$PERF_TOP/conf/cshrc.perf**
 - For sh, ksh, or bash: **.\$PERF_TOP/conf/profile.perf**
3. Verify that there are no errors in the LSF data loader log files.

The data loader log files (*data_loader_name.log.host_name*) are located in the data loader subdirectory of the log directory:

- UNIX: \$PERF_LOGDIR/data_loader
- Windows: %PERF_LOGDIR%\data_loader

Check the status of the Platform Analytics node database connection

1. If you are 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**
2. If you are logged into a UNIX host, source the PERF environment.
 - For csh or tcsh: **source \$PERF_TOP/conf/cshrc.perf**
 - For sh, ksh, or bash: **.\$PERF_TOP/conf/profile.perf**
3. Navigate to the PERF binary directory.
 - UNIX: **cd \$PERF_TOP/version_number/bin**
 - Windows: **cd %PERF_TOP%\version_number\bin**
4. View the status of the node database connection.
 - UNIX: **dbconfig.sh**
 - Windows: **dbconfig**

Check core dump on the Platform Analytics node

Check and enable core dump on the following OS.

Core dump on Linux

1. If you are logged into a UNIX host, source the LSF environment.
 - For csh or tcsh: **source LSF_TOP/conf/cshrc.lsf**
 - For sh or bash: **. LSF_TOP/conf/profile.lsf**
2. If you are logged into a UNIX host, source the PERF environment.
 - For csh or tcsh: **source \$PERF_TOP/conf/cshrc.perf**
 - For sh or bash: **.\$PERF_TOP/conf/profile.perf**

3. Check if core dump is enabled.
 - For `csch` or `tcsh`: **`ulimit -c unlimited`**
 - For `sh` or `bash`: **`ulimit -c`**

If it displays 0, then it is disabled.
4. Enable core dump.
 - For `csch` or `tcsh`: **`limit coredumpsize unlimited`**
 - For `sh` or `bash`: **`ulimit coredump`**
5. Restart the loader controller and apply your changes.

`perfadmin stop all`

`perfadmin start all`
6. Collect the stack trace from the node host.
 - Source the environment variables
 - Use `gdb` to load the core file.

`gdb ${JAVA_HOME}/bin/java core_file`

where *core_file* is the dump core file generated by the Analytics node
 - Print the stack trace: **`bt`**
7. Collect the output from various installations to check if they are correct.

For environment variables: **`env`**

For `csch` or `tcsh`: **`limit`**

For `sh` or `bash`: **`ulimit -a`**

Verify rpm packages that you have installed: **`rpm -qa|grep glibc`**

Core dump on Solaris

1. If you are logged into a UNIX host, source the LSF environment.
 - For `csch` or `tcsh`: **`source LSF_TOP/conf/cshrc.lsf`**
 - For `sh` or `bash`: **`. LSF_TOP/conf/profile.lsf`**
2. If you are logged into a UNIX host, source the PERF environment.
 - For `csch` or `tcsh`: **`source $PERF_TOP/conf/cshrc.perf`**
 - For `sh` or `bash`: **`. $PERF_TOP/conf/profile.perf`**
3. Check if core dump is enabled.
 - For `csch` or `tcsh`: **`ulimit -c unlimited`**
 - For `sh` or `bash`: **`ulimit -c`**

If it displays 0, then it is disabled.
4. Enable core dump.
 - For `csch` or `tcsh`: **`limit coredumpsize unlimited`**
 - For `sh` or `bash`: **`ulimit coredump`**
5. Restart the loader controller and apply your changes.

`perfadmin stop all`

perfadmin start all

6. Collect the stack trace from the node host.

/usr/proc/bin/pstack core_file >pstack.out

/usr/proc/bin/pmap core_file >pmap.out

/usr/proc/bin/pldd core_file >pldd.out

where *core_file* is the dump core file generated by the Analytics node

7. It is recommended that you use dbx to collect stack trace.

- Source the environment variables
- Use dbx to load the core file.

dbx \${JAVA_HOME}/bin/java core_file

- Print the stack trace: **where**

8. Collect the output from various installations to check if they are correct.

For environment variables: **env**

For csh or tcsh: **limit**

For sh or bash: **ulimit -a**

For patches currently installed: **showrev -p**

For detailed information about the packages installed on a system: **pkginfo -l**

Core dump on AIX and HP-UX

1. If you are logged into a UNIX host, source the LSF environment.

- For csh or tcsh: **source LSF_TOP/conf/cshrc.lsf**
- For sh or bash: **. LSF_TOP/conf/profile.lsf**

2. If you are logged into a UNIX host, source the PERF environment.

- For csh or tcsh: **source \$PERF_TOP/conf/cshrc.perf**
- For sh or bash: **.\$PERF_TOP/conf/profile.perf**

3. Check if core dump is enabled.

- For csh or tcsh: **ulimit -c unlimited**
- For sh or bash: **ulimit -c**

If it displays 0, then it is disabled.

4. Enable core dump.

- For csh or tcsh: **limit coredumpsize unlimited**
- For sh or bash: **ulimit coredump**

5. Restart the loader controller and apply your changes.

perfadmin stop all

perfadmin start all

6. It is recommended that you use dbx to collect stack trace.

- Source the environment variables
- Use dbx to load the core file.

```
dbx ${JAVA_HOME}/bin/java core_file
```

where *core_file* is the dump core file generated by the Analytics node

- Print the stack trace: **where**

7. Collect the output from various installations to check if they are correct.

For environment variables: **env**

For csh or tcsh: **limit**

For sh or bash: **ulimit -a**

For release number of the OS: **uname -a**

Debug LSF API

Enable debugging for the LSF API.

1. Set the following environment variables for the current session.

- For sh or bash:

```
export LSF_DEBUG_CMD="LC_EXEC LC_COMM LC_TRACE"
```

```
export LSF_CMD_LOG_MASK=LOG_DEBUG3
```

```
export LSF_CMD_LOGDIR="log_path"
```

```
export LSB_DEBUG_CMD="LC_EXEC LC_COMM LC_TRACE"
```

```
export LSF_CMD_LOG_MASK=LOG_DEBUG3
```

```
export LSF_CMD_LOGDIR="log_path"
```

where *log_path* is the full path where debugging log files are generated.

- For tcsh and tsh: Follow the same commands as sh or bash, but use **setenv** instead of **export**.

2. Restart the loader controller in the same command line session where you set the environment variables.

```
perfadmin stop all
```

```
perfadmin start all
```

3. When data loader start to collect data from LSF, the following log files are generated under the specified directory.

- *lscmd log host_name*

- *bcmd log host_name*

Where *host_name* is the name of the Analytics node host.

Analytics node did not respond

If INFO level messages are not updated for more than one hour in the `ANALYTICS_TOP/log/plc.log.host_name` file, the Analytics node may not respond. Check for the following reasons to resolve this issue.

1. Check if the specified maximum heap size is less than the minimum memory required for the data volume. Check for the following in the log file.

Memory info before gc: *memory in bytes*

Memory info after gc: *memory in bytes*

If the specified heap size is less than the minimum memory requirement, then increase the heap size by changing the java settings in the ANALYTICS_TOP/conf/wsm/wsm_pl c. conf file.

For example: **JAVA_OPTS=-Xms64m -Xmx2048m**

Note:

For Windows 32bit, the maximum heap size that you can set is 1600M.
For Linux / Unix 32bit, you can set it to 4096M. For 64bit machine, you can set it to any value.

2. Check if there is enough disk space for the Analytics node host. If that is the problem, then contact your Administrator to resolve the disk space issue. You need to restart the loader controller once you increase the disk space.

25

Troubleshooting the server

Actions to troubleshoot the Platform Analytics server

- [Check the health of the Platform Analytics server](#) on page 118
- [Check the Platform Analytics server log files](#) on page 118
- [Check the status of the Platform Analytics server database connection](#) on page 119

Check the health of the Platform Analytics server

Use the Platform Analytics Console to verify that the Platform Analytics server is running correctly.

1. Log into the Platform Analytics server host.
2. Launch the Platform Analytics Console.
 - UNIX: **`ANALYTICS_TOP/bin/runconsole.sh`**
 - Windows: Start > Programs > Platform Platform Analytics Server > Platform Platform 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 verify that the scheduled tasks are running correctly according to schedule.

You can also check the data purger scheduled tasks (Part i t i onMai ntenanceGroup*) and compare the data purger settings with your cluster data retention policies.
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.

Check the Platform Analytics server log files

Check the Platform Analytics server log files to verify that there are no errors.

1. Verify that there are no errors in the data purger log file.

The data purger log file (`purger.log. host_name`) is located in the Platform Analytics server log directory:

 - UNIX: `ANALYTICS_TOP/log`
 - Windows: `ANALYTICS_TOP\log`
2. Verify that there are no errors in the event manager log file.

The event manager log file (`eventmanager.log. host_name`) is located in the Platform Analytics server log directory:

 - UNIX: `ANALYTICS_TOP/log`
 - Windows: `ANALYTICS_TOP\log`
3. Verify that there are no errors in the automation manager log file.

The automation manager log file (`automationmanager.log. host_name`) is located in the Platform Analytics server log directory:

 - UNIX: `ANALYTICS_TOP/log`

- Windows: *ANALYTICS_TOP*\log

Check the status of the Platform Analytics server database connection

Use the Platform Analytics Console to verify the Platform Analytics server database connection.

1. Log into the Platform Analytics server host.
2. Launch the Platform Analytics Console.
 - UNIX: ***ANALYTICS_TOP*/bin/runconsole.sh**
 - Windows: Start > Programs > Platform Platform Analytics Server > Platform Platform Analytics Console
3. Click Data Sources in the navigation tree.
4. For each database entry in the main window, test the database connection.
 - a) Right-click the database name and select Edit Data Source.

The Data Source Properties window displays.
 - b) Click Test to test the database connection.

Troubleshooting the server

Customizing Platform Analytics

Platform Analytics customizations allow you to maintain and upgrade your Platform Analytics installation to improve performance and fix issues. Contact Platform Support to find more information on specific customizations to meet your needs or to fix specific issues.

Platform Analytics customizations provided by Platform follow specific conventions. If you create your own customizations, your customizations must follow these same conventions to ensure that your customizations are compatible and are saved if you upgrade your Platform Analytics installation.

Naming conventions

The name of the customization is the same as the package name and identifies the specific customization, allowing Platform Support to easily locate the source code for your specific customization.

The customization name is the module or activity name followed by an underscore (`_`) and a serial number. For Platform customizations, this number is often the support ticket number for the Platform Analytics enhancement or bug fix.

Subdirectories containing files belonging to the customization must have names followed by an underscore and the serial number. Similarly, files belonging to the customization that are located in common directories must also have names followed by an underscore and the serial number.

Node customizations

The following describes conventions and examples of customizations to the Platform Analytics node:

- [Supported files](#) on page 124
- [Customize an existing data loader](#) on page 124
- [Add a new custom data loader](#) on page 125

Supported files

Customizations to the following built-in configuration files (all in the `conf` directory) will remain in the upgraded or patched Analytics node:

- `datasource.xml`
- `log4j.properties`
- `plc.xml`
- `perf.conf`
- All `*.properties` files in the `data loader` subdirectory.
- All `*.xml` files in the `plc` subdirectory.
- `wsm_plc.conf` files in the `wsm` subdirectory.

Customizations to other Platform Analytics files might not remain in an upgrade or patched Platform Analytics node. Therefore, in order to meet Platform Analytics node conventions, customizations to the Platform Analytics node cannot overwrite any Platform Analytics files not in this supported list.

Customize an existing data loader

If you customize an existing data loader, do not directly overwrite the built-in binaries. Instead, you can edit the source code, make file, or `build.xml` file to build binaries with different names by following the naming conventions.

The following describes an example to customize the `lsfpendingreasonloader` to obtain more information for detailed pending reasons:

1. Edit the necessary source code to change or add the necessary required information.
For example, edit the `pendreason.c` file.
2. Edit the make file to build the final `.so` file with a different name (such as appending the serial number).
For example, edit the make file to build the final file named `libpendreason_148781.so`.
3. Change the package name to a different name (such as appending the serial number).
For example, for all files in the `com.platform.perf.data loader.lsf.advanced.pendreason` package, change the package name to `com.platform.perf.data loader.lsf.advanced.pendreason_148781`.
4. Change the Java code to load the new shared library.
For example, in the `com.platform.perf.data loader.lsf.advanced.pendreason_148781.ReadPendReasonJNI.java` file, change the `System.loadLibrary` line to the following:

```
System.loadLibrary("pendreason_148781");
```
5. Edit the `build.xml` file to build the final `.jar` file with a different name.

For example, edit the `build.xml` file to build the `pendreason_148781.jar` file.

6. Copy the existing data loader configuration to a file that follows the customization file naming convention.

For example, copy the existing data loader configuration to `pendingreason_148781.xml`.

7. Edit the new data loader configuration file with the desired attributes.
 - a) Change the `Class` attribute of the `Reader` element to the new class that you specified as the package name.

For example, change the `Class` attribute from

`com.platform.perf.dataloader.lsf.advanced.pendreason` to

`com.platform.perf.dataloader.lsf.advanced.pendreason_148781`.

- b) To add more columns that you want the data loader to collect, edit the `SQL` section.

8. Edit the loader controller configuration file to point to the new data loader configuration file.

For example, the relevant directories and files are as follows:

`ANALYTICS_TOP`

- `conf`
- `dataloader/pendingreason_148781.xml`

The data loader configuration file.

- `plc/plc_lsf_advanced.xml`

The loader controller configuration file related to the pending reason data loader. This file may be modified for the new data loader.

- `lsf/7.0`

Library files collecting Platform LSF 7.0 data.

Similarly, the `ego` directory contains library files collecting EGO-related data, and the `license` directory contains library files collecting license-related data.

- `dataloader/pendingreason_148781.xml`

The data loader configuration file.

- `platform/lib/libpendreason_148781.so`

The shared library file is here.

Add a new custom data loader

Add a new data loader to collect custom data from the cluster.

1. Add the loader controller configuration file for the new data loader to the `ANALYTICS_TOP/conf/plc` directory.

Create a new loader controller configuration file by copying the `plc.xml` file and editing the copied file for your new data loader. It is recommended that you create at least one standalone loader controller configuration file for your custom data loaders.

2. Add the new data loader configuration file to the `ANALYTICS_TOP/conf/dataloader` directory.
3. Add the library files to the corresponding `lib` directories.

For example, to create the License Scheduler workload data loader with serial number 148782, add the following files to the following relevant directories:

ANALYTICS_TOP

- `conf`
- `data_loader/l_s_workload_148782.xml`

The data loader configuration file.

- `data_loader/l_s_workload_148782.properties`

The data loader property file.

- `plc/plc_l_s_workload_148782.xml`

A standalone loader controller configuration file for the new data loader.

- `license/7.0`

Library files collecting Platform LSF License Scheduler 7.0 data.

Similarly, the `ego` directory contains library files collecting EGO-related data, and the `lsf` directory contains library files collecting LSF-related data.

- `lib/l_s_workload_148782.jar`
- `platform/lib/lib_sworkload_148782.so`

The shared library file is here.

Server customizations

The following describes conventions and examples of customizations to the Platform Analytics server:

- [Supported files](#) on page 127
- [Customize an existing Tableau Server workbook](#) on page 127

Supported files

Customizations to the following built-in configuration files (all in the `conf` directory) will remain in the upgraded or patched Analytics server:

- `datasource.xml`
- `log4j.properties`
- `Config.xml`
- `ItemLists.xml`
- `pi.conf`
- All `*.xml` files in the `purger` subdirectory.
- `Package.xml` files in the `packages/workload` subdirectory.

Customizations to other Platform Analytics files might not remain in an upgrade or patched Platform Analytics server. Therefore, in order to meet Platform Analytics server conventions, customizations to the Platform Analytics server cannot overwrite any Platform Analytics files not in this supported list.

Customize an existing Tableau Server workbook

Customizing an existing Tableau Server workbook is not recommended, because the customization is not guaranteed to remain in the upgraded or patched workbook. Instead, copy the existing workbook to a new one following the naming convention. Use the Platform Analytics Designer to customize the new workbook and publish.

Database schema customizations

When customizing the database schema, you should only perform the following actions:

- Create a new object.
- Add a new column to a built-in table.

Do not perform the following actions to customize the database schema:

- Drop a built-in object.
- Rename a built-in object.
- Drop a column from a built-in table.
- Rename a column in a built-in table.
- Replace a built-in view, procedure, package, or trigger.

Built-in objects include tables, views, procedures, packages, indexes, triggers, and sequences.

Customization management

The following describes the conventions when assembling, installing, or viewing the customization packages (or "patches").

- [Assemble the customization package](#) on page 129
- [Install the customization package](#) on page 130
- [View details on the customization packages](#) on page 130

Assemble the customization package

Binary or configuration files in the customization package should keep the same hierarchical structure as it does in the runtime environment. Perform the following to make your customization package compatible with the Platform Analytics patch installer: and add the following text files to this subdirectory:

1. Create a subdirectory named `patch_install` in the top-level directory of your package.
2. Add patch configuration files to the `patch_install` subdirectory.
 - a) Create and add the `patchinfo.txt` file.

Specify a semicolon-separated list containing detailed patch information in the following format:

build_number;build_date;version;dependency;manual_config

where

- *build_number* is the build request number. This is a unique number that distinguishes the patch from other patches. For customizations, specify any unique build number or use a serial number according to the customization naming conventions. For example, 12345.
- *build_date* is the build date in UTC/GMT time in the following numerical format: YYYYMMDDhhmmss. For example, 20111015104104.
- *version* is the version of your Platform Analytics installation. For example, 8. 0.
- *dependency* is the build number of a fix or solution that this patch depends on. If there is more than one fix or solution dependency, separate multiple build numbers with a comma. If there are no dependencies, use `null`. For example, 1234, 2345.
- *manual_configs* specifies if the patch has manual configuration steps before starting the Platform Analytics services. If set to Y, the patch installer does not restart Platform Analytics services after deploying the patch; otherwise, the patch installer will restart the Platform Analytics services after deploying the patch. The default value is N.

For example,

```
12345; 20111015104104; 8. 0; 1234, 2345; Y
```

- b) Create and add the `fixlist.db` file.

Specify a list of bugs fixed in the patch, with each fixed bug on one line in the file. Each line contains the bug tracking number and an optional brief description, ending with a semicolon, as follows:

bug_number[:description];

For example,

```
148781: Added more columns to pendreasonloader;
```

- c) Create and add the `filelist.db` file.

Specify a list of files in your customization. Use a slash (/) in the file paths for both Windows and UNIX.

For example,

```
conf/data/loader/pendreason_148781.xml
conf/plc/plc_lsf_advanced.xml
lsf/7.0/lib/pendreason_148781
lsf/7.0/lib/nux_64-x86/lib/libpendreason_148781.so
```

Install the customization package

1. Navigate to the `ANALYTICS_TOP/patch_tools` directory.
2. Run the Platform Analytics server patch installer.
 - UNIX: **patch_install.sh**
 - Windows: **patch_install.bat**

Note:

- The patch installer will prompt you to specify the patch directory, which is the absolute file path to the extracted directory of your patch.
 - The patch installer will restart the services on the Platform Analytics server.
-

View details on the customization packages

The following commands allow you to view information on the customizations that are applied to the Platform Analytics installation.

- List information on all patches applied to the current Platform Analytics installation directory.
 - UNIX: **pversion.sh -a all**
 - Windows: **pversion.bat -a all**

The latest patch is shown first.

- List information on the last patch that the current file is from.
 - UNIX: **pversion.sh -f *file_name***
 - Windows: **pversion.bat -f *file_name***
- List detailed information on the specified build.
 - UNIX: **pversion.sh -b *build_name***
 - Windows: **pversion.bat -b *build_name***