
Installing Platform Analytics

Platform Analytics
Version 7.6
May 2010



Contents

1	Analytics hosts	3
	System ports	4
	System architecture diagrams	6
2	Licensing	9
3	Plan for the Analytics installation	11
	What you need to do	12
4	Prepare the database for Analytics	17
	What you need to do	18
5	Prepare to install Analytics	23
	What you need to do	24
6	Install the Analytics node	47
	What you need to do	48
7	Install the Analytics server	51
	What you need to do	52
8	Configure the Analytics node	55
	What you need to do	56
9	Verify and configure the Analytics server	61
	What you need to do	62
10	Optimize the Analytics node	65
	What you need to do	66
11	Optimize the Analytics server	71
	What you need to do	72
12	Remove Analytics software from the host	87
	What you need to do	88

Analytics hosts

There are two types of Analytics hosts that you install:

Analytics node	Analytics nodes are hosts that collect data from clusters or license servers. Each node either belongs to a cluster from which Analytics collects data (including license data if connected to a license server), or is a standalone host that collects license data.
Analytics server	The Analytics server manages the data that the Analytics nodes collect. You can perform all server functions using the Analytics Console in the server.

System ports

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

Analytics system ports

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

Analytics server ports

Port name	Default port number	Description	Firewall configuration (all bi-directional)	Configuration file
PIAM_PORT	9091	Internal port for the task scheduler.	N/A	<i>ANALYTICS_TOP/conf/pi.conf</i>
PIEM_PORT	9092	The Analytics event manager uses this port to receive events from the Analytics server and nodes.	Allow connections for this port from each node host to the server host.	<i>ANALYTICS_TOP/conf/pi.conf</i>
Cognos ports	N/A (see)	These ports work internally on the Analytics server host (which is the Cognos server host).	N/A	See Determine the Cognos system ports on page 43.
Web server port	80	End users can use this port to browse cubes and reports.	Allow connections for this port from each browsing host to the server host.	N/A

Analytics node ports

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

Third-party tool ports

Port name	Default port number	Description	Firewall configuration (all bi-directional)
Cognos ports	N/A (see Determine the Cognos system ports on page 43)	These ports work internally on the Analytics server host (which is the Cognos server host).	N/A
Oracle database port	1521	The Analytics server and nodes use this port to communicate with the Oracle database.	Allow connections for this port from the database host to the server and node hosts.
Web server port	80	End users can use this port to browse cubes and reports.	Allow connections for this port from each browsing host to the server host.
License server port	N/A	To collect license usage data from the FLEXnet server, the Analytics node needs to communicate with the FLEXnet server host via TCP.	Allow connections for this port from the node host to each FLEXnet server host.
FLEXnet Manager database port	1521 for Oracle 1433 for MS SQL Server	To collect data from FLEXnet Manager, the Analytics node needs to communicate with the FLEXnet Manager database via TCP.	Allow connections for this port from the node host to the FLEXnet Manager database host.

Cognos system ports

The port numbers for some of the Cognos OLAP tools are randomly assigned during installation. Use the Cognos Configuration Manager to determine the Cognos system ports. For more information, see [Determine the Cognos system ports](#) on page 43.

System architecture diagrams

The following diagrams show architectures of Analytics systems that monitor multiple LSF clusters.

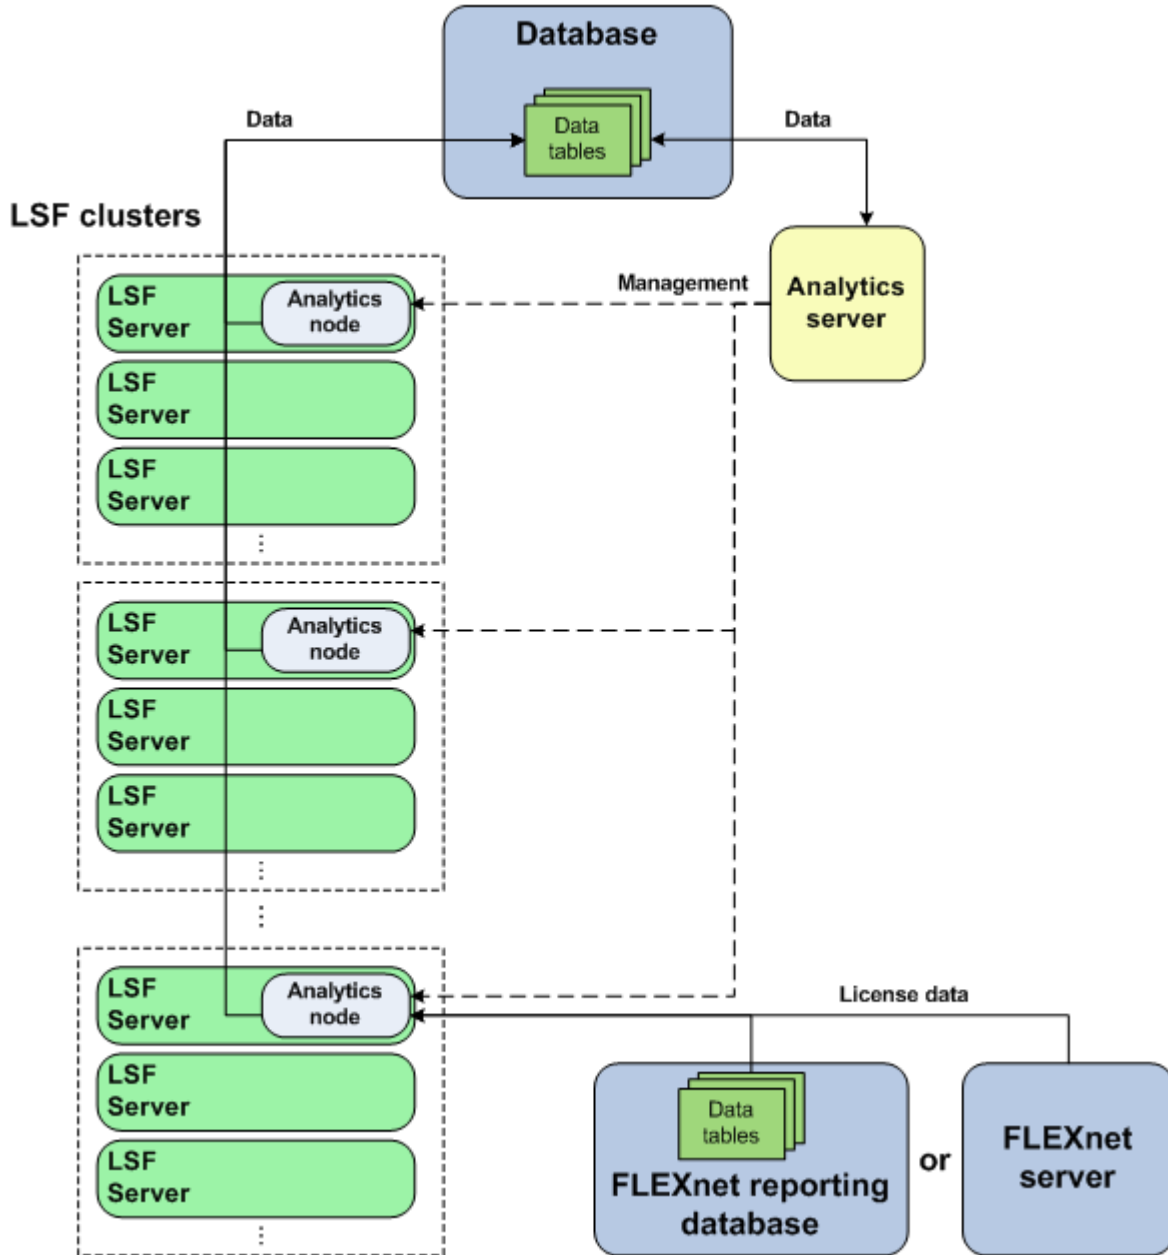


Figure 1: Analytics system architecture where each node is a member of an LSF cluster

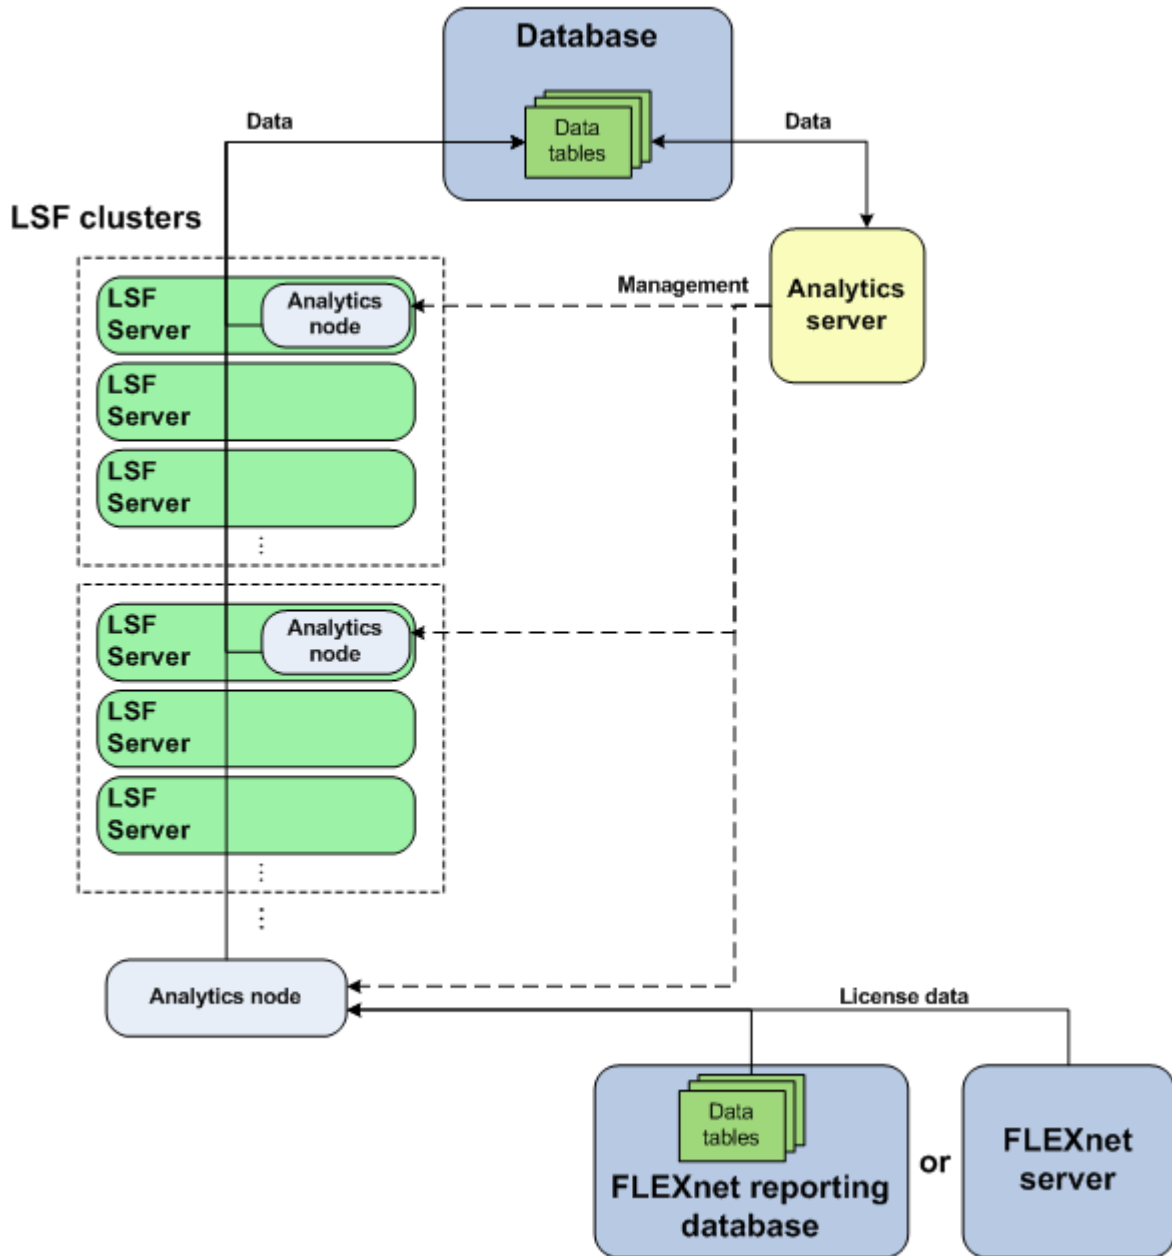


Figure 2: Analytics system architecture with a standalone node connected to a license server

Analytics hosts

Licensing

The Analytics license file includes licenses for data collection and cube building.

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

LSF advanced data collection	The LSF advanced data collection 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	The license data collection license allows you to collect license usage and event data from your FLEXnet servers.
Cube and datamart building	The cube and datamart building license allows you to build datamarts and cubes.

If you have a demo license and obtained a production license, you need to replace the old demo license file in the *PERF_CONF* directory (*\$PERF_CONF* in UNIX and *%PERF_CONF%* in Windows) with the new demo license file. Make sure that both licenses have the same file name (usually *license.dat*).

Plan for the Analytics installation

Select the hosts and download the installation packages for Analytics.

What you need to do

1. [Select the database host](#) on page 12
2. [Select the Analytics server host](#) on page 13
3. [Select the Analytics node hosts](#) on page 14
4. [Obtain the necessary installation files](#) on page 14

After planning for the Analytics installation, prepare the database for Analytics. See [Prepare the database for Analytics](#) on page 17.

Select the database host

If you already have a database host (running Oracle) that you plan to use with Analytics, you need to ensure that Analytics supports your version of the Oracle database, and that there is appropriate space in the database allocated for Analytics. Refer to the *Release Notes for Platform Analytics* for the latest list of supported Oracle database versions.

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

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

Automatic data striping is available in certain RAID devices under software or hardware control. Oracle Automatic Storage Management (ASM) allows ASM files to be either coarse- or fine-striped. You can also achieve data striping with Logical Volume Management (LVM) in Linux.

The following table describes the optimal configuration of the database depending on the size of your cluster:

Cluster size	Database	RAM	CPU	Local hard disk	Table space	Network
Medium (100 - 1000 hosts)	Oracle Enterprise Edition	8GB	4 × 3.0GHz	SCSI 15K RPM RAID 01 or 10 2× RAM size for swap space	100GB for 300 000 finished jobs per day	Gigabit Ethernet
Large (more than 1000 hosts)	Oracle Enterprise Edition with partitioning	16GB	8 × 3.0GHz	SCSI 15K RPM RAID 01 or 10 2× RAM size for swap space	350GB for 1 000 000 finished jobs per day	Gigabit Ethernet

Data striping configuration

When configuring data striping in your RAID 01 or 10 system, the following parameters are important when improving I/O performance:

Stripe width

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

Stripe size

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

You should use a large stripe size of at least 128 KB. The preferred size is as follows:

Stripe size = DB_BLOCK_SIZE × MULTIBLOCK_READ_COUNT

Typically, this is 8 KB × 16 = 128 KB.

You can also use manual data striping. To do this, add multiple files to each tablespace, with each file on a separate disk.

Select the Analytics server host

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

For optimal performance, the Analytics server host should be a dedicated multi-core host with sufficient memory and input/output performance. If the Analytics server is running on a Windows host, you should use the 64-bit version of Windows because Java cannot use more than 1638MB of memory on 32-bit platforms.

Input/output (I/O) performance is an important performance factor in the Analytics server host because cube build performance largely depends on I/O performance. The Analytics server host should have local SCSI disks with RAID 0 to enable concurrent I/O for cube building.

The following table describes the optimal configuration of the Analytics server depending on the size of your cluster:

Cluster size	RAM	CPU	Local hard disk	Network
Medium (100 - 1000 hosts)	4GB	2 × 3.0GHz	Either one of: <ul style="list-style-type: none"> • SATA 7200 RPM • SCSI 10K RPM • Solid state disk 60 GB free space and 2× RAM size for swap space	Gigabit Ethernet
Large (more than 1000 hosts)	16GB	8 × 3.0GHz	Either one of: <ul style="list-style-type: none"> • SATA 10K RPM • SCSI 15K RPM • Solid state disk 200 GB free space and 2× RAM size for swap space	Gigabit Ethernet

Note:

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

Select the Analytics node hosts

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

For optimal performance of your Analytics node, the host should be running on a high-bandwidth network. Since network bandwidth is an important performance bottleneck for the Analytics nodes, the Analytics node host should have a Gigabit Ethernet connection with the database host. If the Analytics node is running on a Windows host, you should use the 64-bit version of Windows because Java cannot use more than 1638MB of memory on 32-bit platforms.

The following table describes the optimal configuration of the Analytics node depending on the size of the clusters in which the node resides:

Cluster size	RAM	CPU	Local hard disk	Network
Medium (100 - 1000 hosts)	2GB	3.0GHz	SATA 7200 RPM 10GB of free space 2x RAM size for swap space	Gigabit Ethernet
Large (more than 1000 hosts)	4GB	2 x 3.0GHz	SATA 10K RPM 10GB of free space 2x RAM size for swap space	Gigabit Ethernet

Obtain the necessary installation files

1. Obtain the necessary files for installing Platform Analytics.

You need the following files to install Platform Analytics:

- Platform Analytics server installation package
- Platform Analytics node installation package
- Platform Analytics data schema package
- Platform Analytics documentation package
- Platform Analytics license file
- Oracle JDBC driver (version 10.2.0.3 only). This driver is available from the following URL:

http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html

Attention:

You must specifically use the 10.2.0.3 version of the Oracle JDBC driver, because other versions do not work correctly with Analytics.

2. Obtain the necessary files for installing the Cognos OLAP tools.

You need installation packages for the following Cognos OLAP tools:

- Sun Java Directory Server, found in the Cognos Supplementary Software package

The installation file or directory name contains "sun_one" or "sol supp".

- Cognos PowerPlay Enterprise Server

The installation file or directory name contains "ppes".

- Cognos PowerPlay Transformer Edition

The installation file or directory name contains "ppt w".

3. Obtain the necessary files for installing the latest version of Adobe Acrobat Reader.

You can download the latest version from the Adobe web site (<http://www.adobe.com>), or use the installer included with the Cognos Supplemental Software package.

Prepare the database for Analytics

Create tablespaces and data schema, then configure the Oracle database to prepare for the Analytics installation.

What you need to do

Check the following:

- The Oracle database host is properly configured and running.
- You have a user name, password, and URL to access the database.
- The Oracle database and the database host meet the Analytics system requirements.

See [Plan for the Analytics installation](#) on page 11.

1. [Create the Oracle tablespaces for the database schema](#) on page 18
2. [Create an Oracle database schema](#) on page 19
3. [Optimize the Oracle database](#) on page 20

After preparing the database for Analytics, prepare the Analytics server host. See [Prepare to install Analytics](#) on page 23.

Create the Oracle tablespaces for the database schema

Contact your Oracle database administrator (DBA) to create the tablespaces according to these specifications.

Create tablespaces to prepare the Oracle database for the Analytics database schema. You may create partitioned or non-partitioned data schema.

Partitioned data schema maximizes Analytics performance, reduces database administration complexity, and allows the data loaders and data transformers to work more efficiently with the data volume at a larger cluster scale.

For optimal database performance, we recommend that you use partitioned data schema and create seven tablespaces for Analytics. If you are using Oracle Standard Edition, you can only use non-partitioned data schema. If you are using non-partitioned data schema you only need to create two tablespaces.

1. Determine the scale of your cluster, based on the number of hosts, and the corresponding initial tablespace size.

Refer to the following table to determine the scale of your cluster and the initial tablespace.

Number of hosts in the cluster	Scale of the cluster	Initial tablespace size in total
Less than 100	Tiny	4 GB (4096 MB)
100 - 1000	Small	20 GB (20480 MB)
1000 - 3000	Medium	200 GB (204800 MB)
Greater than 3000	Large	500 GB (512000 MB)

2. Create tablespaces with the "AUTOEXTEND" and "EXTENT MANAGEMENT LOCAL" options with the initial tablespace size according to the scale of your cluster.

Note:

AUTOEXTEND specifies that the tablespaces are automatically extended with a specified size for the NEXT parameter. Specifying the NEXT parameter assures that data files grow with consistent, reusable extend sizes.

EXTENT MANAGEMENT LOCAL allows you to specify a locally-managed tablespace. You should generally not specify a dictionary-managed tablespace.

```
CREATE TABLESPACE tablespace_name DATAFILE 'data_file_path' SIZE tablespace_sizeM
AUTOEXTEND ON 2048M EXTENT MANAGEMENT LOCAL;
```

where

- *tablespace_name* is the name of the tablespace
- *data_file_path* is the file path to the . ora Oracle data file.
- *tablespace_size* is the initial size of the tablespace according to the scale of your cluster, in MB. For example, for small clusters (100 to 1000 hosts), this is 20480.

For example, for small clusters (100 to 1000 hosts),

```
CREATE TABLESPACE TS_DATA_01 DATAFILE '/oradata/ts_data_01_01.ora' SIZE 20480M
AUTOEXTEND ON 2048M EXTENT MANAGEMENT LOCAL;
```

and repeat this command for the remaining tablespaces (for example, tablespaces TS_DATA_02 to TS_DATA_07 for partitioned data schema and tablespace TS_DATA_02 for non-partitioned data schema).

Create an Oracle database schema

Check the following:

- The Analytics database user must have the following:
 - Roles: Connect and resource.
 - System privileges: Create view.
1. Log into any host with access to the database host as the Analytics database user.
 2. Extract the database schema package.
 3. Navigate to the directory containing your database schema scripts.
 - If you are creating a partitioned data schema, navigate to `schema/Oracle/partitioned`
You must be using Oracle Enterprise Edition to work with partitioned databases.
 - If you are creating a non-partitioned data schema, navigate to `schema/Oracle/nonpartitioned`
 4. Use `sqlplus` to run the `grant_privilege.sql` script with DBA privileges to grant additional, required privileges to the Analytics database user.

```
sqlplus "sys/password@connect_string as sysdba" @grant_privilege.sql user_name
```

where

- *password* is the password for the `sys` user on the database.
- *connect_string* is the named SQLNet connection for this database (optional).
- *user_name* is the name of the Analytics database user.

5. Use sql plus to run the `create_pa76_schema.sql` script to create a database schema.

- To create a partitioned data schema:

```
sqlplus user_name/password@connect_string @create_pa76_schema.sql data_tablespace
index_tablespace data_tablespace,index_tablespace,partitioned_tablespace_list cluster_scale
```

- To create a non-partitioned data schema:

```
sqlplus user_name/password@connect_string @create_pa76_schema.sql data_tablespace
index_tablespace
```

where

- *user_name* is the user name on the database.
- *password* is the password for this user name on the database.
- *connect_string* is the named SQLNet connection for this database (optional).
- *data_tablespace* is the name of the tablespace where you intend to store the table schema.
- *index_tablespace* is the name of the tablespace where you intend to store the index.
- *partitioned_tablespace_list* is a comma-separated list of the remaining tablespaces where you intend to store the partitioned tables and indexes. Spaces are not allowed in this list.
- *cluster_scale* is a letter representing the scale of your cluster (T - Tiny, S - Small, M - Medium, L - Large).

For example, if you have a new small cluster (100 to 1000 hosts),

If you created seven tablespaces for a partitioned data schema,

```
sqlplus Analytics_owner/mypasswd @create_pa76_schema.sql
TS_DATA_01 TS_DATA_02 TS_DATA_01,TS_DATA_02,TS_DATA_03,TS_DATA_04,TS_DATA_
05,TS_DATA_06,TS_DATA_07 S
```

If you created two tablespaces for a non-partitioned data schema,

```
sqlplus Analytics_owner/mypasswd @create_pa76_schema.sql TS_DATA_01 TS_DATA_02
```

6. Partitioned databases only. Wait for thirty minutes after you created the database schema, then verify that the partitions for all tables have been created correctly.

Use a SQL query to check the `sys_partition_events` table for any partitioning error messages.

```
SELECT * FROM sys_partition_events
```

If there is an error message due to insufficient tablespace size, you need to expand the tablespace size and create the partitions again by running the following command on the SQL Plus command line:

```
EXEC CREATE_PARTITIONS;
```

Optimize the Oracle database

Optional. Modify the Oracle database for optimal performance with Analytics.

1. Modify the Oracle database initialization parameters.

The following table describes the parameters you need to modify:

Parameter name	Value
sga_target (Oracle 10g or newer only)	Total memory × (35 to 40%)

Parameter name	Value
pga_aggregate_target	Total memory × (35 to 40%)
open_cursors	300
processes	300
db_block_size	8192
db_file_multiblock_read_count	16
optimizer_index_cost_adj	30
undo_retention	28800

If you are using Oracle 9i, you also need to modify the following parameters:

Parameter name	Value
log_buffer	1048576
sga_max_size	Total memory × (35 to 40%)
large_pool_size	104857600
shared_pool_size	314572800
java_pool_size	0
db_cache_size	sga_max_size - large_pool_size - shared_pool_size - java_pool_size - log_buffer
sort_area_size	1048576

2. Increase the size of the redo log files to decrease log file synchronization time.

Increase the number of redo log groups to four and increase the size of each of the redo log groups to 1.5 GB.

In addition, you should redistribute these redo log files to a physical hard disk other than the one containing the operating system and the tablespaces for Analytics data.

3. Increase the UNDO and TEMP tablespaces.

Change the size of the UNDO and TEMP tablespaces to 20 GB, or as large as possible.

For example, if your cluster has a large volume of records (such as a full workload scale LSF cluster consisting of 5000 hosts and 3 million finished jobs a day with 40 000 license usage combinations), you may need to set the TEMP and UNDO tablespace size to 40 GB or more.

4. Change the default Oracle statistics gathering job (GATHER_STATS_JOB) to weekly base.

To change the job to weekly base, run the following PL/SQL block in SQL*Plus with DBA privileges:

```
BEGIN
  BEGIN
    DBMS_SCHEDULER.DROP_JOB(JOB_NAME => 'GATHER_STATS_JOB', FORCE => TRUE);
  EXCEPTION
    WHEN OTHERS THEN
      NULL;
  END;
  DBMS_SCHEDULER.CREATE_JOB(JOB_NAME          => 'GATHER_STATS_JOB',
    PROGRAM_NAME    => 'GATHER_STATS_PROG',
    JOB_CLASS       => 'AUTO_TASKS_JOB_CLASS',
```

Prepare the database for Analytics

```
START_DATE      => SYSTIMESTAMP,  
REPEAT_INTERVAL => 'freq=daily; byday=SAT; byhour=0; byminute=0; bysecond=0',  
END_DATE        => NULL,  
ENABLED         => TRUE,  
COMMENTS        => 'Oracle defined automatic optimizer statistics collection job.');
```

END;
/

Prepare to install Analytics

Prepare the server host for the Analytics installation.

What you need to do

Check the following:

- The Analytics server host meets the detailed system and software requirements.

For more details, refer to [Plan for the Analytics installation](#) on page 11.

- The Analytics server host meets the software requirements for installing the Cognos tools.

The Analytics server host must be a Windows or Solaris host to install the Cognos tools. For additional requirements, refer to the Cognos documentation.

If you do not wish to generate cubes with the Analytics server, you may proceed with the installation if the host does not meet the Cognos requirements. The Analytics server host can be a UNIX or Windows host if you are not installing Cognos tools.

- The required ports for the Analytics and database hosts are free. If a firewall exists, all these ports must be open in the firewall. All the required ports are TCP.

For more details, refer to [Analytics hosts](#) on page 3.

1. [Install the Cognos tools](#) on page 24
 - [Install the Cognos 7.4 MR3 tools](#) on page 25
 - [Install the Cognos 7.4 MR4 tools](#) on page 26
2. [Configure the Cognos tools](#) on page 27
 - [Configure the Cognos 7.4 MR3 tools on a Solaris host](#) on page 28
 - [Configure the Cognos 7.4 MR3 tools on a Windows host](#) on page 32
 - [Configure the Cognos 7.4 MR4 tools on a Solaris host](#) on page 35
 - [Configure the Cognos 7.4 MR4 tools on a Windows host](#) on page 39
3. [Change the default publishing mode of cube reports from PDF to HTML](#) on page 42
4. [Test the Cognos OLAP tools](#) on page 42
5. [Determine the Cognos system ports](#) on page 43
6. [Enable LSF job event data exporting \(LSF 6.2, 7.0.4, 7.0.5, and 7.0.6 only\)](#) on page 44

After preparing the Analytics server, install the Analytics node package on a host in the cluster.

Install the Cognos tools

Check the following:

- You have a properly configured web server using the fully qualified domain name of the host, and Analytics supports this configuration.
- The character set of your host is set to English.
- If you are installing Cognos in a Windows host, check the following:

- If the host is in a domain, the host name must be fully qualified.

To verify that your Windows host name is fully qualified, select System Properties from the Control Panel and click Computer Name. The Full computer name field must indicate a fully-qualified host name including the domain name.

- The primary DNS suffix of the host is set properly.

To check the primary DNS suffix, select System Properties from the Control Panel and click Computer Name. Click Change, then More. The DNS suffix should be the same as the domain name.

- The web server must be CGI-enabled in order to log into the Cognos server.

This document assumes that you are using the Apache web server for Solaris hosts and Internet Information Services (IIS) for Windows hosts.

To verify that your Windows hosts (running IIS) are CGI-enabled, check that All Unknown CGI Extensions is set to Allowed.

You can skip this step if you do not wish to generate cubes with the Analytics server.

Refer to the Cognos documentation for further details on how to install the Cognos OLAP tools. You must install either Cognos 7.4 MR3 or MR4:

- [Install the Cognos 7.4 MR3 tools](#) on page 25
- [Install the Cognos 7.4 MR4 tools](#) on page 26

Install the Cognos 7.4 MR3 tools

1. Log into the Analytics server host.
 - In Solaris, log in as root.
 - In Windows, log in as the intended Analytics administrator account.
2. Install the Sun Java Directory Server using the typical installation mode.

The directory server on the Analytics server stores configuration information.

Remember the following when installing the Sun Java Directory Server:

- The installation file or directory name contains "sun_one" or "sol supp".
 - Select Install Sun Java System Directory Server instead of Install and Configure Sun Java System Directory Server, if the two options exist.
 - Select the Sun Java System Servers installation option and all Directory Server components for installation.
 - Take note of the Directory Manager DN password after you specify it during installation. You will need this password later when you configure the Cognos tools.
 - Take note of the server port number that you specified. You will need this port number later to access the Sun Java System Server Console.
3. Verify that the directory server processes are running.
 - a) Run the Sun Java System Server Console.
 - Solaris: `/var/Sun/mps/startconsole &`
 - Windows: Select Start > Programs > Sun Java(TM) System Server Products > Sun Java(TM) System *version_number*.

Note:

You may need to change the port number to the server port number that you noted earlier. If you forget this port number, edit the `adm.conf` file to find it. This file is located in the `admin_serv/conf` subdirectory of the Sun Java Directory Server installation directory.

- b) In the navigation tree, navigate to *domain_name*>*server_host_name*>Server Group>Directory Server.
 - c) Check that the status of the server is "Started".
4. Install Cognos PowerPlay Enterprise Server using the default installation mode.
- Cognos PowerPlay Enterprise Server is required for publishing your cubes.
- Remember the following when installing the Cognos PowerPlay Enterprise Server:
- The installation file or directory name contains "ppes".
 - After completing the installation, be sure to select Exit the Installation Wizard without configuring components.
5. Install Cognos PowerPlay Transformer Edition using the default installation mode.
- Cognos PowerPlay Transformer Edition is an OLAP modeling tool that transforms source data into multidimensional models used to create cubes.
- Remember the following when installing Cognos PowerPlay Transformer Edition.
- The installation file or directory name contains "pptw".
 - You must choose the same installation directory as Cognos PowerPlay Enterprise Server.
 - After completing the installation, be sure to select Exit the Installation Wizard without configuring components.
6. Install Adobe Acrobat Reader using the default installation mode.
- The Adobe Acrobat Reader is required to open the Cognos PowerPlay reports. You can download the latest version, or use the installer included with the Cognos Supplemental Software package.

Install the Cognos 7.4 MR4 tools

1. Log into the Analytics server host.
 - In Solaris, log in as root.
 - In Windows, log in as the intended Analytics administrator account.
2. Install the Sun Java Directory Server from the IBM Cognos Supplementary Software CD (or the equivalent network directory) using the typical installation mode.

The directory server on the Analytics server stores configuration information.

- a) Run the installation program and follow the dialog prompts to complete the installation.
 - In Solaris, navigate to the `supp_sol` installation directory and run `./install` to launch the installation program.
 - In Windows, run **Setup.exe**.

Select Install and Configure Sun Java System Directory Server and accept the default values.

If you already have Cognos installed on your Solaris host, you may see a warning that Setup has detected some running processes. If so, you should exit Setup, launch the Configuration Manager, and stop the Cognos services before running the installation program again.

- b) After the installer finishes copying the files, select the appropriate configuration option.
 - If you installed on a Windows host, in the Component Configuration dialog, select the Configure components with defaults and click Next.
The Windows configuration may take a while to complete.
 - If you installed on a Solaris host, select Start Configuration Manager and click Finish.

The Solaris configuration is not complete until you run the Configuration Manager and successfully apply the changes. If you exit without applying the changes, you must start the Configuration Manager from within the installation directory. Navigate to the `bin` subdirectory of the Cognos installation directory and run `configure_dsi` to start the Configuration Manager.

- c) If you installed on a Solaris host, change the port number and the full computer name using the Configuration Manager.

For example, to set the port to 389, the SSL port to 636 on the `hostA` host in the `test.example.com` domain, run the following commands:

```
cd hostA.3rd\Party\Sun\Java\System\Server\Products\Installation.General.Access\
Manager\Directory\Server\Information.
set Port.=389
set Directory\Server\SSL\Port.=636
cd /
cd hostA.3rd\Party\Sun\Java\System\Server\Products\Installation.General.Sun\Java\
System\Information.
set Full\Computer\Name.= hostA.test.example.com
cd /
apply
```

The `apply` command starts the configuration and activation of the Solaris directory server. The configuration may take a while to complete.

- d) Close the Configuration Manager to complete the installation process.

`exit`

3. Install Cognos PowerPlay Enterprise Server from the IBM Cognos PowerPlay Enterprise Server CD (or the equivalent network directory) using the default installation mode.

Cognos PowerPlay Enterprise Server is required for publishing your cubes.

- In Solaris, navigate to the `ppes_sol` installation directory and run `./install` to launch the installation program.
- In Windows, run `Setup.exe`.

After completing the installation, be sure to exit the installation wizard without configuring components.

4. Install Cognos PowerPlay Transformer Edition from the IBM Cognos Transformer Edition CD (or the equivalent network directory) using the default installation mode.

Cognos PowerPlay Transformer Edition is an OLAP modeling tool that transforms source data into multidimensional models used to create cubes.

- In Solaris, navigate to the `pptr_sol` installation directory and run `./install` to launch the installation program.
- In Windows, run `Setup.exe`.

After completing the installation, be sure to exit the installation wizard without configuring components.

5. Install Adobe Acrobat Reader using the default installation mode.

The Adobe Acrobat Reader is required to open the Cognos PowerPlay reports. You can download the latest version, or use the installer included with the Cognos Supplemental Software package.

Configure the Cognos tools

You must have installed the Cognos tools (Sun Java Directory Server, Cognos PowerPlay Enterprise Server, and Cognos PowerPlay Transformer Edition).

Skip this step if you did not install the Cognos tools.

- [Configure the Cognos 7.4 MR3 tools on a Solaris host](#) on page 28
- [Configure the Cognos 7.4 MR3 tools on a Windows host](#) on page 32
- [Configure the Cognos 7.4 MR4 tools on a Solaris host](#) on page 35
- [Configure the Cognos 7.4 MR4 tools on a Windows host](#) on page 39

Configure the Cognos 7.4 MR3 tools on a Solaris host

Check the following:

- You installed the Cognos 7.4 MR3 tools (Sun Java Directory Server, Cognos PowerPlay Enterprise Server, and Cognos PowerPlay Transformer Edition).
- Your host has access to the X-Windows environment:
 - If you are connected to the intended Analytics server host locally, the host must be running X-Windows.
 - If you are connected to the intended Analytics server host via `telnet`, you must be running `xserver` on your local host.

Refer to the Cognos documentation for further details on how to configure the Cognos OLAP tools.

You need to configure Cognos to use the fully qualified domain name for the host (the Analytics default) rather than the short host name (the Cognos default).

1. Log into the Analytics server host as root.
2. If you connected to the intended Analytics server host via `telnet` and are running `xserver` on a local host, set your display environment.

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

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

- For `csh` or `tcsh`:
setenv DISPLAY hostname:0.0
- For `sh`, `ksh`, or `bash`:
DISPLAY=hostname:0.0
export DISPLAY

3. Create a directory server instance in the Sun Java Directory Server.
 - a) Navigate to the Sun Java Directory Server installation directory.

cd /var/Sun/mps

- b) Start the Sun Java System Server Console.

./startconsole &

When prompted, specify the user name, password, and the default port number of your Directory Server.

- c) In the navigation tree, navigate to `domain_name > server_host_name > Server Group`.
 - d) Right-click `Server Group` and select `Create Instance Of > Directory Server`.
 - e) Take the default settings and click `OK` to create the directory server instance.
4. Start the Cognos Configuration Manager (`configure`) from the `cer5/bin` subdirectory of the Cognos installation directory.

If you installed Cognos in the default directory (/usr/cognos), run **/usr/cognos/er5/bin/configure**.

5. Configure the directory server general settings with the base distinguished name (DN) from the Sun Java directory server.

- a) Navigate to *server_host_name* > Services > Access Manager - Directory Server > General
- b) Set the parameter to configure the directory server.

set "Are you sure you want to configure this directory server?"=Yes

- c) Set the base distinguished name (DN) to the same value as the Sun Java Directory Server.

For example, if your Analytics server host is hostA in the domain example.com domain,

set "Base distinguished name (DN)"=dc=domain,dc=example,dc=com

Note:

To find the value in the Sun Java Directory Server, log into the Sun Java System Server Console and select **Servers and Applications > domain_name > User Directory subtree**.

- d) View and verify the parameter settings, especially the DN and port settings.

Is

- e) Optional. Change the unrestricted user password to the same value as the Directory Manager DN password that you specified when installing the Sun Java Directory Server.

The password parameter is an encrypted value, so you need to use the `encrypt` command to change the value.

encrypt "Unrestricted User password"

Follow the subsequent prompts to enter the new password.

- f) Validate and apply the new parameter settings.

validate

apply

6. Configure the directory server settings with the base distinguished name (DN) from the Sun Java Directory Server.

- a) Navigate to *server_host_name* > Services > Access Manager - Runtime > Authentication Source > Directory Server
- b) Set the base distinguished name (DN) to the same value as the Sun Java Directory Server.

For example, if your Analytics server host is hostA in the domain example.com domain,

set "Base distinguished name (DN)"=dc=domain,dc=example,dc=com

Note:

To find the value in the Sun Java directory server, log into the server interface and select **Servers and Applications > domain_name > User Directory subtree**.

- c) Validate and apply the new parameter settings.

validate

apply

- d) Navigate to *server_host_name*.
- e) Start the Cognos service.

start

- f) Verify that the Cognos service is running.

ps -ef|grep cognos

- 7. Configure Cognos Upfront with the fully qualified domain name of the Analytics server host.

- a) Navigate to *server_host_name* > Cognos Shared > Runtime Parameters > Server Configuration > Upfront > Upfront Server Group.
- b) View the current parameter settings, especially the gateway parameter.

ls

- c) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain *example.com* domain, and the gateway parameter currently ends with */cognos/cgi-bin/upfcgi.exe*,

set gateway=http://hostA.domain.example.com/cognos/cgi-bin/upfcgi.exe

- d) Validate and apply the new parameter settings.

validate

apply

- 8. Configure Cognos PowerPlay Enterprise Server with the fully qualified domain name of the Analytics server host.

- a) Navigate to *server_host_name* > Cognos Shared > Runtime Parameters > Server Configuration > Cognos PowerPlay Enterprise Server > PPES Server Group.
- b) View the current parameter settings, especially the gateway parameter.

ls

- c) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain *example.com* domain, and the gateway parameter currently ends with */cognos/cgi-bin/ppdscgi.exe*,

set gateway=http://hostA.domain.example.com/cognos/cgi-bin/ppdscgi.exe

- d) Validate and apply the new parameter settings.

validate

apply

- 9. Configure Cognos Access Manager with the fully qualified domain name of the Analytics server host.

- a) Navigate to *server_host_name* > Cognos Shared > Runtime Parameters > Server Configuration > Access Manager > AccManLogon.
- b) View the current parameter settings, especially the gateway parameter.

ls

- c) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain *example.com* domain, and the gateway parameter currently ends with */cognos/cgi-bin/login.exe*,

set gateway=http://hostA.domain.example.com/cognos/cgi-bin/login.exe

- d) Navigate to *server_host_name* > Cognos Shared > Runtime Parameters > Server Configuration > Access Manager > AccManAdmin.
- e) View the current parameter settings, especially the gateway parameter.

ls

- f) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is `hostA` in the domain `example.com`, and the gateway parameter currently ends with `/cognos/cgi-bin/accessadmin.exe`,

set gateway=http://hostA.domain.example.com/cognos/cgi-bin/accessadmin.exe

- g) Validate and apply the new parameter settings.

validate

apply

10. Configure Cognos Upfront File Manager with the fully qualified domain name of the Analytics server host.

- a) Navigate to `server_host_name > Cognos Shared > Runtime Parameters > Server Configuration > Upfront File Manager > Upfront File Manager Server Group`.

- b) View the current parameter settings, especially the gateway parameter.

Is

- c) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is `hostA` in the domain `example.com`, and the gateway parameter currently ends with `/cognos/cgi-bin/fmcgi.exe`,

set gateway=http://hostA.domain.example.com/cognos/cgi-bin/fmcgi.exe

- d) Validate and apply the new parameter settings.

validate

apply

11. Validate and apply the new configuration to the server.

- a) Navigate to `server_host_name`.

- b) Validate and apply the new parameter settings.

validate

apply

- c) Restart the Cognos service.

stop

start

12. Close the Cognos Configuration Manager.

13. Configure the web server.

- a) Edit the Apache web server configuration file.

Edit `/usr/apache/conf/httpd.conf`

- b) Configure aliases for your web server to include the `webcontent`, `cgi-bin`, and `documentation` subdirectories in the `cer5` subdirectory.

Add `/cognos/cgi-bin` as a `ScriptAlias` to `cgi-bin`, `/cognos` as an `Alias` to `webcontent`, and `/cognos/help` as an `Alias` to `documentation`.

For example, if you installed Cognos to the default directory (`/usr/cognos`), add the following lines to `httpd.conf`:

```
ScriptAlias /cognos/cgi-bin "/usr/cognos/cer5/cgi-bin"
<Directory "/usr/cognos/cer5/cgi-bin">
AllowOverride None
Options None
```

```
</Directory>
Alias /cognos "/usr/cognos/cer5/webcontent"
<Directory "/usr/cognos/cer5/webcontent">
Options None
AllowOverride None
Order Allow,Deny
Allow from All
</Directory>
Alias /cognos/help "/usr/cognos/cer5/documentati on"
<Directory "/usr/cognos/cer5/documentati on">
Options None
AllowOverride None
Order Allow,Deny
Allow from All
</Directory>
```

- c) Verify that you have execute permissions in the webcontent, cgi-bin, and documentati on subdirectories in the cer5 subdirectory.
- d) Restart the web server.

```
cd /usr/apache/bin
```

```
./apachectl stop
```

```
./apachectl start
```

- e) Verify that the web server is running.

```
ps -ef|grep apache
```

14. Secure your Cognos PowerPlay Enterprise Server.

- a) Run the PowerPlay Server administration script (ppsrvadm. sh) from the cer5/bin subdirectory of the Cognos installation directory.

If you installed Cognos to the default directory (/usr/cognos), run **/usr/cognos/cer5/bin/ppsrvadmin.sh**.

Note:

The first time you run the script, you will be prompted to select the server name and to enter your user name and password. The default user name is **administrator** with a blank password, and the server name is the full domain name of your intended Analytics server.

- b) Select the option to secure your server.
- c) Log into the Server Administrator interface and start the service.

You can access the Cognos server interface by entering the /cognos subdirectory of your Analytics server domain name in your browser. Log in using the administrator account with a blank password.

For example, if your Analytics server host is hostA in the domain. example. com domain, enter **http://hostA.domain.example.com/cognos** in your browser. If this does not work, try using the web server port number. For the default port of 80, try **http://hostA.domain.example.com:80/cognos**.

Configure the Cognos 7.4 MR3 tools on a Windows host

You must have installed the Cognos 7.4 MR4 tools (Sun Java Directory Server, Cognos PowerPlay Enterprise Server, and Cognos PowerPlay Transformer Edition) prior to following these steps.

Refer to the Cognos documentation for further details on how to configure the Cognos OLAP tools.

You need to configure Cognos to use the fully qualified domain name for the host (the Analytics default) rather than the short host name (the Cognos default).

1. Log into the Analytics server host as the intended Analytics administrator account.

2. Create a directory server instance in the Sun Java Directory Server.
 - a) Start the Sun Java System Server Console.
Select Start > Programs > Sun Java(TM) System Server Products > Sun Java(TM) System *version_number*.
When prompted, specify the user name, password, and the default port number of your Directory Server.
 - b) In the navigation tree, navigate to *domain_name* > *server_host_name* > Server Group.
 - c) Right-click Server Group and select Create Instance Of > Directory Server.
 - d) Take the default settings and click OK to create the directory server instance.
3. Start the Cognos Configuration Manager.
 - a) Click Start > Programs > Cognos *version_number* > Tools > Configuration Manager.
 - b) Select Open the current configuration.
4. Stop all Cognos processes.
In the navigation tree, right-click your Analytics server host name and select Stop.
5. Configure the directory server general settings with the base distinguished name (DN) from the Sun Java directory server.
 - a) Navigate to *server_host_name* > Services > Access Manager - Directory Server > General
 - b) Set the Are you sure you want to configure this directory server? parameter to Yes.
 - c) Set the Base distinguished name (DN) parameter to the same value as the Sun Java directory server.

For example, if your Analytics server host is hostA in the domain. example. com domain, set the parameter to dc=domain,dc=example,dc=com

Note:

To find the value in the Sun Java directory server, log into the Sun Java server interface and select **Servers and Applications** > ***domain_name*** > **User Directory subtree**.

- d) Verify the other parameter settings, especially the DN and port settings.
- e) Validate the new parameter settings.
In the navigation tree, right-click General and select Validate Selection.
- f) Apply the new parameter settings.

In the navigation tree, right-click General and select Apply Selection.

6. Configure the directory server settings with the base distinguished name (DN) from the Sun Java directory server.
 - a) Navigate to *server_host_name* > Services > Access Manager - Runtime > Authentication Source > Directory Server
 - b) Set the Base distinguished name (DN) parameter to the same value as the Sun Java directory server.

For example, if your Analytics server host is hostA in the domain. example. com domain, set the parameter to dc=domain,dc=example,dc=com

Note:

To find the value in the Sun Java directory server, log into the Sun Java server interface and select **Servers and Applications** > ***domain_name*** > **User Directory subtree**.

- c) Verify the other parameter settings, especially the DN and port settings.
- d) Optional. Set the Unrestricted User password parameter to the same value as the Directory Manager DN password that you specified when installing the Sun Java Directory Server.
- e) Validate the new parameter settings.

In the navigation tree, right-click Directory Server and select Validate Selection.

- f) Apply the new parameter settings.

In the navigation tree, right-click Directory Server and select Apply Selection.

7. Start the Cognos processes.

In the navigation tree, right-click your Analytics server host name and select Start.

8. Configure Cognos server groups with the fully qualified domain name of the Analytics server host.

- a) Below the navigation tree, click the Server Configuration tab.
- b) In the navigation tree, click All Server Groups.
- c) For each Server Group under All Server Groups, change the Gateway URL and Web Server parameters to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain. example.com domain, change all instances of **http://hostA** to **http://hostA.domain.example.com**.

- d) Below the navigation tree, click the Components tab.
- e) Validate the new parameter settings.

In the navigation tree, right-click your Analytics server host name and select Validate Selection.

- f) Apply the new parameter settings.

In the navigation tree, right-click your Analytics server host name and select Apply Selection.

9. Restart the Cognos processes.

- a) In the navigation tree, right-click your Analytics server host name and select Stop.
- b) In the navigation tree, right-click your Analytics server host name and select Start.

10. Close the Cognos Configuration Manager.

11. Configure aliases for your web server.

You need to share and map the webcontent, cgi-bin, and Documentat i on directories to the web server.

- a) Navigate to the cer5\bin subdirectory of the Cognos installation directory.

If you installed Cognos to the default directory, navigate to C:\Program Files\Cognos\cer5\bin

- b) Right-click the webcontent directory and select Sharing.
- c) Click the Web Sharing tab and specify the following fields:

- Share this folder: Enabled
- Alias: cognos

- d) Click Edit Properties and specify the following fields:

- Access permissions: Read
- Application permissions: Scripts

- e) Right-click the cgi-bin directory and select Sharing.

- f) Click the Web Sharing tab and specify the following fields:

- Share this folder: Enabled

- Alias: `cognos\cgi - bin`
- g) Click Edit Properties and specify the following fields:
 - Access permissions: Read
 - Application permissions: Execute (includes scripts)
- h) Right-click the `document at i on` directory and select Sharing.
- i) Click the Web Sharing tab and specify the following fields:
 - Share this folder: Enabled
 - Alias: `cognos\hel p`
- j) Click Edit Properties and specify the following fields:
 - Access permissions: Read
 - Application permissions: Scripts
- k) Verify that all accounts (users in the "Everyone" group) have execute permissions in the `webcontent`, `cgi - bin`, and `document at i on` directories.

Important:

If you are installing on a Windows 2003 host, you will not be able to view the dynamic content from Cognos. You must also follow the steps in <http://support.microsoft.com/kb/315122> to view the dynamic content.

12. Secure your Cognos PowerPlay Enterprise Server.

- a) Run the PowerPlay Server administration program (`PPSrvAdm.exe`) from the `cer5\bin` subdirectory of the Cognos installation directory.

If you installed Cognos to the default directory, run "**C:\Program Files\Cognos\cer5\bin\PPSrvAdm.exe**".

Note:

The first time you run the program, you will be prompted to select the server name and to enter your user name and password. The default user name is **administrator** with a blank password, and the server name is the full domain name of your intended Analytics server.

- b) Select the option to secure your server.

You can access the Cognos server interface by entering the `/cognos` subdirectory of your Analytics server domain name in your browser. Log in using the administrator account with a blank password.

For example, if your Analytics server host is `hostA` in the `example.com` domain, enter **`http://hostA.domain.example.com/cognos`** in your browser. If this does not work, try using the web server port number. For the default port of 80, try **`http://hostA.domain.example.com:80/cognos`**.

Configure the Cognos 7.4 MR4 tools on a Solaris host

Check the following:

- You installed the Cognos 7.4 MR4 tools (Sun Java Directory Server, Cognos PowerPlay Enterprise Server, and Cognos PowerPlay Transformer Edition).
- Your host has access to the X-Windows environment:
 - If you are connected to the intended Analytics server host locally, the host must be running X-Windows.

- If you are connected to the intended Analytics server host via telnet, you must be running xserver on your local host.

Refer to the Cognos documentation for further details on how to configure the Cognos OLAP tools.

You need to configure Cognos to use the fully qualified domain name for the host (the Analytics default) rather than the short host name (the Cognos default).

1. Log into the Analytics server host as root.
2. If you connected to the intended Analytics server host via telnet and are running xserver on a local host, set your display environment.

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

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

- For csh or tcsh:

```
setenv DISPLAY hostname:0.0
```

- For sh, ksh, or bash:

```
DISPLAY=hostname:0.0
```

```
export DISPLAY
```

3. Start the Cognos Configuration Manager (configure) from the cer5/bin subdirectory of the Cognos installation directory.

If you installed Cognos in the default directory (/usr/cognos), run **/usr/cognos/cer5/bin/configure**.

4. Configure the directory server general settings with the base distinguished name (DN) from the Sun Java directory server.
 - a) Navigate to *server_host_name* > Services > Access Manager - Directory Server > General
 - b) Set the parameter to configure the directory server.

```
set "Are you sure you want to configure this directory server?"=Yes
```

- c) View and verify the parameter settings, especially the DN and port settings.

```
ls
```

- d) Optional. Change the unrestricted user password to the same value as the Directory Manager DN password that you specified when installing the Sun Java Directory Server.

The password parameter is an encrypted value, so you need to use the encrypt command to change the value.

```
encrypt "Unrestricted User password"
```

Follow the subsequent prompts to enter the new password.

- e) Validate and apply the new parameter settings.

```
validate
```

```
apply
```

5. Configure the directory server settings with the base distinguished name (DN) from the Sun Java Directory Server.
 - a) Navigate to *server_host_name* > Services > Access Manager - Runtime > Authentication Source > Directory Server

- b) Validate and apply the new parameter settings.
 - validate**
 - apply**
 - c) Navigate to *server_host_name*.
 - d) Start the Cognos service.
 - start**
 - e) Verify that the Cognos service is running.
 - ps -ef|grep cognos**
6. Configure Cognos Upfront with the fully qualified domain name of the Analytics server host.
- a) Navigate to *server_host_name* > Cognos Shared > Runtime Parameters > Server Configuration > Upfront > Upfront Server Group.
 - b) View the current parameter settings, especially the gateway parameter.
 - ls**
 - c) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain. example. com domain, and the gateway parameter currently ends with /cognos/cgi-bin/upfcgi.exe,

 - set gateway=http://hostA.domain.example.com/cognos/cgi-bin/upfcgi.exe**
 - d) Validate and apply the new parameter settings.
 - validate**
 - apply**
7. Configure Cognos PowerPlay Enterprise Server with the fully qualified domain name of the Analytics server host.
- a) Navigate to *server_host_name* > Cognos Shared > Runtime Parameters > Server Configuration > Cognos PowerPlay Enterprise Server > PPES Server Group.
 - b) View the current parameter settings, especially the gateway parameter.
 - ls**
 - c) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain. example. com domain, and the gateway parameter currently ends with /cognos/cgi-bin/ppdscgi.exe,

 - set gateway=http://hostA.domain.example.com/cognos/cgi-bin/ppdscgi.exe**
 - d) Validate and apply the new parameter settings.
 - validate**
 - apply**
8. Configure Cognos Access Manager with the fully qualified domain name of the Analytics server host.
- a) Navigate to *server_host_name* > Cognos Shared > Runtime Parameters > Server Configuration > Access Manager > AccManLogon.
 - b) View the current parameter settings, especially the gateway parameter.
 - ls**
 - c) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain. example. com domain, and the gateway parameter currently ends with /cognos/cgi-bin/lugin.exe,

set gateway=http://hostA.domain.example.com/cognos/cgi-bin/login.exe

- d) Navigate to *server_host_name* > Cognos Shared > Runtime Parameters > Server Configuration > Access Manager > AccManAdmin.
- e) View the current parameter settings, especially the gateway parameter.

Is

- f) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain example.com domain, and the gateway parameter currently ends with /cognos/cgi-bin/accessadmin.exe,

set gateway=http://hostA.domain.example.com/cognos/cgi-bin/accessadmin.exe

- g) Validate and apply the new parameter settings.

validate

apply

- 9. Configure Cognos Upfront File Manager with the fully qualified domain name of the Analytics server host.

- a) Navigate to *server_host_name* > Cognos Shared > Runtime Parameters > Server Configuration > Upfront File Manager > Upfront File Manager Server Group.
- b) View the current parameter settings, especially the gateway parameter.

Is

- c) Change the gateway parameter to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain example.com domain, and the gateway parameter currently ends with /cognos/cgi-bin/fmcgi.exe,

set gateway=http://hostA.domain.example.com/cognos/cgi-bin/fmcgi.exe

- d) Validate and apply the new parameter settings.

validate

apply

- 10. Validate and apply the new configuration to the server.

- a) Navigate to *server_host_name*.
- b) Validate and apply the new parameter settings.

validate

apply

- c) Restart the Cognos service.

stop

start

- 11. Close the Cognos Configuration Manager.

- 12. Configure the web server.

- a) Edit the Apache web server configuration file.

Edit /usr/apache/conf/httpd.conf

- b) Configure aliases for your web server to include the webcontent, cgi-bin, and documentat ion subdirectories in the cer5 subdirectory.

Add /cognos/cgi-bin as a ScriptAlias to cgi-bin, /cognos as an Alias to webcontent, and /cognos/help as an Alias to documentat ion.

For example, if you installed Cognos to the default directory (`/usr/cognos`), add the following lines to `httpd.conf`:

```
ScriptAlias /cognos/cgi-bin "/usr/cognos/cer5/cgi-bin"
<Directory "/usr/cognos/cer5/cgi-bin">
AllowOverride None
Options None
</Directory>
Alias /cognos "/usr/cognos/cer5/webcontent"
<Directory "/usr/cognos/cer5/webcontent">
Options None
AllowOverride None
Order Allow,Deny
Allow from All
</Directory>
Alias /cognos/help "/usr/cognos/cer5/documentation"
<Directory "/usr/cognos/cer5/documentation">
Options None
AllowOverride None
Order Allow,Deny
Allow from All
</Directory>
```

- c) Verify that you have execute permissions in the `webcontent`, `cgi-bin`, and `documentation` subdirectories in the `cer5` subdirectory.
- d) Restart the web server.

```
cd /usr/apache/bin
```

```
./apachectl stop
```

```
./apachectl start
```

- e) Verify that the web server is running.

```
ps -ef|grep apache
```

13. Secure your Cognos PowerPlay Enterprise Server.

- a) Run the PowerPlay Server administration script (`ppsrvadm.sh`) from the `cer5/bin` subdirectory of the Cognos installation directory.

If you installed Cognos to the default directory (`/usr/cognos`), run `/usr/cognos/cer5/bin/ppsrvadmin.sh`.

Note:

The first time you run the script, you will be prompted to select the server name and to enter your user name and password. The default user name is **administrator** with a blank password, and the server name is the full domain name of your intended Analytics server.

- b) Select the option to secure your server.
- c) Log into the Server Administrator interface and start the service.

You can access the Cognos server interface by entering the `/cognos` subdirectory of your Analytics server domain name in your browser. Log in using the administrator account with a blank password.

For example, if your Analytics server host is `hostA` in the `domain.example.com` domain, enter `http://hostA.domain.example.com/cognos` in your browser. If this does not work, try using the web server port number. For the default port of 80, try `http://hostA.domain.example.com:80/cognos`.

Configure the Cognos 7.4 MR4 tools on a Windows host

You must have installed the Cognos 7.4 MR4 tools (Sun Java Directory Server, Cognos PowerPlay Enterprise Server, and Cognos PowerPlay Transformer Edition) prior to following these steps.

Refer to the Cognos documentation for further details on how to configure the Cognos OLAP tools.

You need to configure Cognos to use the fully qualified domain name for the host (the Analytics default) rather than the short host name (the Cognos default).

1. Log into the Analytics server host as the intended Analytics administrator account.
2. Start the Cognos Configuration Manager.
 - a) Click Start > Programs > Cognos *version_number* > Tools > Configuration Manager.
 - b) Select Open the current configuration.
3. Stop all Cognos processes.

In the navigation tree, right-click your Analytics server host name and select Stop.

4. Configure the directory server general settings with the base distinguished name (DN) from the Sun Java directory server.
 - a) Navigate to *server_host_name* > Services > Access Manager - Directory Server > General
 - b) Set the Are you sure you want to configure this directory server? parameter to Yes.
 - c) Verify the other parameter settings, especially the DN and port settings.
 - d) Validate the new parameter settings.

In the navigation tree, right-click General and select Validate Selection.

- e) Apply the new parameter settings.

In the navigation tree, right-click General and select Apply Selection.

5. Configure the directory server settings with the base distinguished name (DN) from the Sun Java directory server.
 - a) Navigate to *server_host_name* > Services > Access Manager - Runtime > Authentication Source > Directory Server
 - b) Verify the other parameter settings, especially the DN and port settings.
 - c) Optional. Set the Unrestricted User password parameter to the same value as the Directory Manager DN password that you specified when installing the Sun Java Directory Server.
 - d) Validate the new parameter settings.

In the navigation tree, right-click Directory Server and select Validate Selection.

- e) Apply the new parameter settings.

In the navigation tree, right-click Directory Server and select Apply Selection.

6. Start the Cognos processes.

In the navigation tree, right-click your Analytics server host name and select Start.
7. Configure Cognos server groups with the fully qualified domain name of the Analytics server host.
 - a) Below the navigation tree, click the Server Configuration tab.
 - b) In the navigation tree, click All Server Groups.
 - c) For each Server Group under All Server Groups, change the Gateway URL and Web Server parameters to use the fully qualified domain name.

For example, if your Analytics server host is hostA in the domain. example.com domain, change all instances of **http://hostA** to **http://hostA.domain.example.com**.

- d) Below the navigation tree, click the Components tab.
- e) Validate the new parameter settings.

In the navigation tree, right-click your Analytics server host name and select Validate Selection.

- f) Apply the new parameter settings.

In the navigation tree, right-click your Analytics server host name and select Apply Selection.

8. Restart the Cognos processes.
 - a) In the navigation tree, right-click your Analytics server host name and select Stop.
 - b) In the navigation tree, right-click your Analytics server host name and select Start.
9. Close the Cognos Configuration Manager.
10. Configure aliases for your web server.

You need to share and map the `webcontent`, `cgi-bin`, and `Documentation` directories to the web server.

- a) Navigate to the `cer5\bin` subdirectory of the Cognos installation directory.

If you installed Cognos to the default directory, navigate to `C:\Program Files\Cognos\cer5\bin`
- b) Right-click the `webcontent` directory and select Sharing.
- c) Click the Web Sharing tab and specify the following fields:
 - Share this folder: Enabled
 - Alias: `cognos`
- d) Click Edit Properties and specify the following fields:
 - Access permissions: Read
 - Application permissions: Scripts
- e) Right-click the `cgi-bin` directory and select Sharing.
- f) Click the Web Sharing tab and specify the following fields:
 - Share this folder: Enabled
 - Alias: `cognos\cgi-bin`
- g) Click Edit Properties and specify the following fields:
 - Access permissions: Read
 - Application permissions: Execute (includes scripts)
- h) Right-click the `Documentation` directory and select Sharing.
- i) Click the Web Sharing tab and specify the following fields:
 - Share this folder: Enabled
 - Alias: `cognos\help`
- j) Click Edit Properties and specify the following fields:
 - Access permissions: Read
 - Application permissions: Scripts
- k) Verify that all accounts (users in the "Everyone" group) have execute permissions in the `webcontent`, `cgi-bin`, and `documentation` directories.

Important:

If you are installing on a Windows 2003 host, you will not be able to view the dynamic content from Cognos. You must also follow the steps in <http://support.microsoft.com/kb/315122> to view the dynamic content.

11. Secure your Cognos PowerPlay Enterprise Server.
 - a) Run the PowerPlay Server administration program (`PPSrvAdm.exe`) from the `cer5\bin` subdirectory of the Cognos installation directory.

If you installed Cognos to the default directory, run "**C:\Program Files\Cognos\cer5\bin\PPSrvAdm.exe**".

Tip:

The first time you run the program, you will be prompted to select the server name and to enter your user name and password. The default user name is **administrator** with a blank password, and the server name is the full domain name of your intended Analytics server.

- b) Select the option to secure your server.

You can access the Cognos server interface by entering the /cognos subdirectory of your Analytics server domain name in your browser. Log in using the administrator account with a blank password.

For example, if your Analytics server host is hostA in the domain example.com domain, enter **http://hostA.domain.example.com/cognos** in your browser. If this does not work, try using the web server port number. For the default port of 80, try **http://hostA.domain.example.com:80/cognos**.

Change the default publishing mode of cube reports from PDF to HTML

In Cognos 7.4, the default publishing mode of cube reports changed from HTML to PDF. To change the default publishing mode back to HTML, you need to change the PPSroot.cfx configuration file before installing the Analytics server.

1. Log into the Analytics server host.
 - In Solaris, log in as root.
 - In Windows, log in as the intended Analytics administrator account.
2. Edit the PPSroot.cfx configuration file in the PowerPlay Enterprise Server directory.
 - In Solaris, if you installed Cognos to the default directory (/usr/cognos), edit /user/cognos/cer5/ppserver/PPSroot.cfx
 - If you installed Cognos to the default directory ("C:\Program Files\Cognos"), edit "C:\Program Files\Cognos\cer5\ppserver\PPSroot.cfx".

3. Search for the UpfrontDefaultLaunchMode parameter.

For example, this parameter will appear as follows:

```
UpfrontDefaultLaunchMode="TX, R"
```

The "R" represents the PDF default launch mode.

4. Change the UpfrontDefaultLaunchMode parameter to HTML.

For example, change this parameter to the following:

```
UpfrontDefaultLaunchMode="TX, Q"
```

The "Q" represents the HTML default launch mode.

5. Save the PPSroot.cfx file to save your changes.

If you publish a cube report, then open Cognos Upfront, you will open the report in HTML format.

Test the Cognos OLAP tools

Verify that the Cognos OLAP tools are properly installed and configured.

1. Log into the Analytics server host.
 - In Solaris, log in as root.
 - In Windows, log in as the intended Analytics administrator account.
2. Test the Cognos PowerPlay Enterprise Server.
 - a) Log into the host.
 - b) Run the PowerPlay Server administration script.
 - In Solaris, run **`./ppsrvadmin.sh`**
 Run the PowerPlay Server administration script (`ppsrvadmin.sh`) from the `cer5/bin` subdirectory of the Cognos installation directory.

 If you installed Cognos to the default directory (`/usr/cognos`), run **`/usr/cognos/cer5/bin/ppsrvadm.sh`**.
 - In Windows, run the PowerPlay Server administration program (`PPSrvAdm.exe`) from the `cer5\bin` subdirectory of the Cognos installation directory.

 If you installed Cognos to the default directory ("`C:\Program Files\Cognos`"), run **`"C:\Program Files\Cognos\cer5\bin\PPSrvAdm.exe"`**.
 - c) Select a sample cube.
 - d) Publish the cube by selecting **Tools > Publish Report Links to Upfront** or **Tools > Publish to portal**.
3. Open a cube in Cognos Upfront.

You can use Upfront to view the cube you published from your PowerPlay Enterprise Server. Cubes and reports published from the PowerPlay Enterprise Server are displayed as NewsItems in Upfront.

 - a) Launch the Upfront gateway from your web browser.

 For example, if your Analytics server host is `hostA` in the `domain.example.com` domain, open the following URL:

`http://hostA.domain.example.com/cognos`
 - b) If prompted, log on using the default user ID (**Administrator**) with a blank password.

 The NewsIndex opens with the sample cube you just published from the PowerPlay Enterprise Server.
 - c) Click the hypertext link (with the name of the sample cube) next to the cube icon.

 The cube opens in the PowerPlay Web Explorer.

Determine the Cognos system ports

The port numbers for some of the Cognos OLAP tools are randomly assigned during installation. If you need to ensure that the Cognos ports are open, you must determine what port numbers are assigned to the Cognos OLAP tools.

1. Log into the Analytics server host.
2. Start the Cognos Configuration Manager.
 - Windows: Click **Start > Programs > Cognos *version_number* > Tools > Configuration Manager** and select **Open the current configuration**.
 - Solaris: Start the Cognos Configuration Manager (`configure`) from the `cer5/bin` subdirectory of the Cognos installation directory.

If you installed Cognos in the default directory (`/usr/cognos`), run `/usr/cognos/cer5/bin/configure`.

3. Determine the Cognos PowerPlay Enterprise Server port number.

- a) Navigate to `server_host_name > Cognos PowerPlay Enterprise Server > Dispatcher > General`.

The Cognos PowerPlay Enterprise Server port number is shown in the PowerPlayer Server Port Number parameter.

4. Determine the Cognos Upfront port numbers.

- a) Navigate to `server_host_name > Upfront > Data Store > General`.

- The Cognos Upfront data store port number is shown in the Data Store Port parameter.
- The Cognos Upfront search engine port number is shown in the Search Engine Port parameter.

- b) Navigate to `server_host_name > Upfront > Server > Administration Connection`.

The Cognos Upfront server administration port number is shown in the Server Administration Port parameter.

5. Determine the Access Manager Server port number.

- a) Navigate to `server_host_name > Services > Access Manager - Server > General > Ticket Service`.

The Access Manager Server port number is shown in the Port parameter.

6. Determine the Access Manager Directory Server port numbers.

- a) Navigate to `server_host_name > Services > Access Manager - Directory Server > General`.

- The Access Manager Directory Server port number is shown in the Port parameter. The default value is 389, which is the same as the Sun Java Directory Server port number.
- The Access Manager primary ticket service port number is part of the Primary ticket service parameter, which is in the following format:

```
server_host_name:port_number
```

Enable LSF job event data exporting (LSF 6.2, 7.0.4, 7.0.5, and 7.0.6 only)

Your cluster must be running one of the following:

- LSF 6.2 with the latest Maintenance Pack
- LSF 7.0.4
- LSF 7.0.5
- LSF 7.0.6

By default, LSF 6.2, 7.0.4, 7.0.5, and 7.0.6 do not enable the `lsb.stream` file for the exporting of LSF job event data.

If you want the Analytics node to collect LSF cluster data from your LSF 6.2, 7.0.4, 7.0.5, or 7.0.6 cluster, you need to enable the `lsb.stream` file because Analytics requires this file for the data loaders to obtain job data.

1. Log into a host in the LSF cluster.
2. Edit the `lsb.params` file.

- UNIX: `$LSF_ENVDIR/lsbatch/cluster_name/configdir/lsb.params`
 - Windows: `%LSF_ENVDIR%\lsbatch\cluster_name\configdir\lsb.params`
3. In the `lsb.params` file, edit the Parameters section to enable the exporting of LSF job event data to the `lsb.stream` file.

Add the following lines to the Parameters section:

```
# Enable streaming of lsbatch system events
ENABLE_EVENT_STREAM=y
# Determines the location of the lsb.stream file. This parameter is optional.
# The default location is: $LSB_SHAREDIR/{clustername}/logdir/stream.
# EVENT_STREAM_FILE=/tmp/lsb.mystream
# Determines the maximum size of the lsb.stream file. This parameter is optional.
# The default size is 100MB.
# MAX_EVENT_STREAM_SIZE=10000
```

4. Reconfigure `mbatchd` to apply these changes.

badmin mbdrestart

5. To verify that these changes are in effect, verify that the `lsb.stream` files exist.

By default, `lsb.stream` is located at the following directories:

- UNIX: `$LSB_SHAREDIR/cluster_name/logdir/stream`
- Windows: `%LSB_SHAREDIR%\cluster_name\logdir\stream`

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

Prepare to install Analytics

Install the Analytics node

Install the Analytics node package on a UNIX or Windows host. You need to have one Analytics node host in each cluster that Analytics will handle.

What you need to do

Check the following:

- The Analytics node host meets the detailed system and software requirements.
For more details, refer to [Plan for the Analytics installation](#) on page 11.
- The Analytics node host operating system can handle daylight savings time correctly.
- If you want the Analytics node to collect LSF cluster data, check the following:
 - You have access to the primary LSF administrator account.
 - The Analytics node host must be an LSF server in the cluster with access to the LSF event file (`lsb.stream`).
- If you want the Analytics node to collect LSF cluster data from an LSF 6.2 cluster, you must install the latest Maintenance Pack on your LSF 6.2 cluster.
- If you are installing on a UNIX host, your host must have access to the X-Windows environment:
 - If you are connected to the intended Analytics node host locally, the host must be running X-Windows.
 - If you are connected to the intended Analytics node host via `telnet`, you must be running `xserver` on your local host.
- The required port for the Analytics hosts is free. If a firewall exists, this TCP port must be open in the firewall.

For more details, refer to [Analytics hosts](#) on page 3.

- The Oracle database host is properly configured and running:
 - You have a user name, password, and URL to access the database.
 - There is appropriate space in the database allocated for Analytics.
 - Analytics supports your version of the Oracle database.
- The Analytics node host has access to the Oracle JDBC driver (version 10.2.0.3 only). This driver is available from the following URL:

http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html

- The Analytics node host has access to the Analytics license file (`license.dat`).
1. Log into the intended Analytics node host.
 - In UNIX, if you want to update the `rc` boot scripts, log in as root.
 - In Windows, log in as the intended Analytics administrator account. If you want the Analytics node to collect LSF cluster data, you must log in as the primary LSF administrator.
 2. If you connected to the UNIX host via `telnet` and are running `xserver` on a local host, set your display environment.

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

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

- For `csch` or `tcsh`:
setenv DISPLAY hostname:0.0
- For `sh`, `ksh`, or `bash`:
DISPLAY=hostname:0.0

export DISPLAY

where *hostname* is your local host.

3. If you are installing on a UNIX host and you want the Analytics node to collect LSF cluster data, source the LSF environment.
 - For `csh` or `tcsh`: **source *LSF_TOP/conf/cshrc.lsf***
 - For `sh`, `ksh`, or `bash`: **. *LSF_TOP/conf/profile.lsf***
4. Run the Analytics node installation package.
 - UNIX: ***janalytics76_node_platform.bin***
 - Windows: ***analytics76_node_platform.exe***
5. Follow the dialog prompts to complete the installation.

To install the Analytics node, you may need to input the following:

- Analytics license file location
- Data collection settings
- Installation directory

The default Analytics node installation directory is as follows:

- Windows: `C:\PlatformAnalytics76_node`
- UNIX: `/opt/platformanalytics76_node`

In this documentation, *ANALYTICS_TOP* refers to the top-level Analytics installation directory in your host.

- LSF cluster settings (if you enabled LSF data collection)

By default, this is the value of the `LSF_ENVDIR` environment variable, which is the location of `lsf.conf`.

- PLC communication port

You need to specify a port for the loader controller (`plc`), which uses 4046 by default.

- Analytics administrator details

Specify the user name and group of the intended Analytics administrator account, which must already exist. If you want the Analytics node to collect LSF cluster data, you must specify the primary LSF administrator account.

- Data source (database host) properties.

The Oracle user name and password is generally the account that you used to create your data schema.

Replace the sample values in the URL field with the values for your database. Ask your Oracle database administrator for the SID and the port number.

- Set RC properties, if you installed on a UNIX host as `root`.

Enable this setting to update the `rc.boot` script.

6. If you are installing on a UNIX host, source the PERF environment.
 - For `csh` or `tcsh`: **source *ANALYTICS_TOP/conf/cshrc.perf***
 - For `sh`, `ksh`, or `bash`: **. *ANALYTICS_TOP/conf/profile.perf***
7. Verify that the `plc` service has started.

Run **`perfadmin list`** and verify that `plc` is running.

Install the Analytics node

8. If the `plc` service has not already started, manually start the `plc` service.

perfadmin start plc

9. Verify that the data loaders are running.

- Windows: **plcclient -s**
- UNIX: **plcclient.sh -s**

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

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

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

After installing the Analytics node package, install the Analytics server package.

Install the Analytics server

Install the Analytics server on a UNIX or Windows host.

What you need to do

Check the following:

- The Analytics server host meets the detailed system and software requirements.
For more details, refer to [Plan for the Analytics installation](#) on page 11.
- The Analytics server host operating system can handle daylight savings time correctly.
- You removed any previous installations of Analytics from your host.
- The intended Analytics administrator account must have privileges to remove the `lsb.stream` account from the `LSF_TOP/work/cluster_name/logdir/stream` directory.

Typically, this account would have the same privileges as the LSF administrator account.

- If you are installing on a UNIX host, your host must have access to the X-Window environment:
 - If you are connected to the intended Analytics server host locally, the host must be running X-Window.
 - If you are connected to the intended Analytics server host via `telnet`, you must be running `xserver` on your local host.
- The required ports for the Analytics and database hosts are free. If a firewall exists, all these ports must be open in the firewall. All the required ports are TCP.

For more details, refer to [Analytics hosts](#) on page 3.

- The Oracle database host is properly configured and running:
 - You have a user name, password, and URL to access the database server.
 - There is appropriate space in the database allocated for Analytics.
 - Analytics supports your version of the Oracle database.
- The Analytics server host has access to the Oracle JDBC driver (version 10.2.0.3 only). This driver is available from the following URL:

http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html

- The Analytics server host has access to the Analytics license file (`license.dat`).

1. Log into the intended Analytics server host.

- In UNIX, if you log in as root, you will be able to update the `rc.boot` scripts and specify the user name and group of the Analytics administrator during the installation.

Otherwise, log in as the intended Analytics administrator account.

- In Windows, log in as the intended Analytics administrator account.

2. If you connected to the UNIX host via `telnet` and are running `xserver` on a local host, set your display environment.

Test your display by running `xclock` or another X-Window application.

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

- For `csch` or `tcsh`:
setenv DISPLAY hostname:0.0
- For `sh`, `ksh`, or `bash`:
DISPLAY=hostname:0.0

export DISPLAY

where *hostname* is your local host.

3. Run the Analytics server installation package.
 - UNIX: ***JAnalytics76_server_platform.bin***
 - Windows: ***Analytics76_server_platform.exe***
4. Follow the dialog prompts to complete the installation.

To install the Analytics server, you may need to input the following:

- Installation directory
 - Windows: C: \Anal yti cs76
 - UNIX: /opt/Anal yti cs76

In this documentation, *ANALYTICS_TOP* refers to the top-level Analytics installation directory in your host.

- Analytics license file location
- Communication ports

The Analytics server requires the use of two consecutive ports, starting with the base port. The default base port is 9090, therefore, the Analytics server uses ports 9090 and 9091.

- Email notification settings

You can receive notifications of events via email. This means you will be made aware of potential problems in your cluster without having to constantly monitor the Analytics Console.

Note:

You should enable the email notification during the initial installation rather than at a later time. If you do not enable email notification, you will need to check the Analytics Console regularly to make sure that Analytics is working properly.

- Analytics administrator details

If you logged in as root on a UNIX host, specify the user name and group of the intended Analytics administrator account, which must already exist.

If you logged into a Windows host, or as a non-root account on a UNIX host, your current account will be the Analytics administrator.

- Database connection properties

The user name and password is generally the account that you used to create your data schema.

- Database partitioning
- Cognos server properties
- Set RC properties, if you installed on a UNIX host as root.

Enable this setting to update the *rc.boot* script.

5. If you are installing on a UNIX (Solaris or Linux) host, start the Analytics server.

Run ***daemons.sh start*** from the *ANALYTICS_TOP/bin* directory.

After installing the Analytics server package, configure the Analytics node.

Install the Analytics server

Configure the Analytics node

Configure the Analytics node to work with your cluster to work together efficiently.

What you need to do

1. [Stop the redundant PERF services \(in LSF 7 clusters only\)](#) on page 56
2. [Configure the LSF cluster for the host core utilization data loader](#) on page 56
3. [Enable FLEXnet data loaders in your cluster](#) on page 57
4. [Enable FLEXnet Manager data loaders in your cluster](#) on page 58
5. [Install and configure the Platform FLEXnet Manager scripts](#) on page 59

After configuring the Analytics node, configure the Analytics server.

Stop the redundant PERF services (in LSF 7 clusters only)

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

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

1. Log into any host in the LSF cluster as the cluster administrator.
2. If you logged into a UNIX host, source the LSF environment.
 - For `cs`h or `tc`sh: **source `LSF_TOP/conf/cshrc.lsf`**
 - For `sh`, `ksh`, or `bash`: **. `LSF_TOP/conf/profile.lsf`**
3. Navigate to the PERF binary directory.
 - UNIX: **cd `$PERF_TOP/version_number/bin`**
 - Windows: **cd `%PERF_TOP%\version_number\bin`**
4. Stop the loader controller (`plc`) and data purger (`purger`) services.
perfadmin stop plc
perfadmin stop purger
5. Rename the loader controller configuration file (`plc.xml`) to another file extension to remove it from the list of LSF services.
For example, rename `plc.xml` to `plc.xml.bak`.
6. Rename the data purger configuration file (`purger.xml`) to another file extension to remove it from the list of LSF services.
For example, rename `purger.xml` to `purger.xml.bak`.
7. Restart EGO on the LSF master host to activate these changes.

egosh ego restart `lsf_master_host_name`

Configure the LSF cluster for the host core utilization data loader

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

1. Log into the Analytics node host.
2. If you logged into a UNIX host, source the LSF environment.
 - For `csh` or `tcsh`: **source `LSF_TOP/conf/cshrc.lsf`**
 - For `sh`, `ksh`, or `bash`: **. `LSF_TOP/conf/profile.lsf`**
3. Copy the `elim.coreutil` execution files from the Analytics directory to the relevant LSF directory.
 - UNIX:


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

where **`os_type`** is `aix`, `hpux`, `linux`, or `solaris`, depending on the specific UNIX operating system.
 - Windows (LSF 7.0.x only):


```
copy ANALYTICS_TOP\elim\windows\elim.coreutil.exe LSF_TOP\7.0\etc
```
 - Windows (LSF 6.2 only):


```
copy ANALYTICS_TOP\elim\windows\elim.coreutil.exe LSF_TOP\etc
```
4. Edit the `LSF_TOP/conf/lsf.shared` file and add `CORE_UTIL` as a resource.

Add the following text to the Resource section:

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

For example,

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

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

Add the following text to the Resource section:

```
CORE_UTIL [default]
```

For example,

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

6. Reconfigure the LSF cluster to apply your changes.

```
badadmin reconfig
```

```
lsadmin reconfig
```

Enable FLEXnet data loaders in your cluster

Check the following:

- You configured Analytics to collect FLEXnet license data.
 - The FLEXnet usage data loader and the FLEXnet license server are using the same time zone.
1. Log into the Analytics node host.
 2. Modify the FLEXnet usage data loader configuration file to set up a server list pointing to the FLEXnet license servers.
 - a) Edit the FLEXnet usage data loader configuration file.

Configure the Analytics node

- UNIX: `ANALYTICS_TOP/conf/dataloader/flexlicusage.properties`
 - Windows: `ANALYTICS_TOP\conf\dataloader\flexlicusage.properties`
 - b) Change the `ServerList` (or `FileName`) parameter to point to the FLEXnet license servers.
3. Modify the FLEXnet events data loader configuration file to point to the license log file.
 - a) Edit the FLEXnet events data loader configuration file.
 - UNIX: `ANALYTICS_TOP/conf/dataloader/flexlicevents.properties`
 - Windows: `ANALYTICS_TOP\conf\dataloader\flexlicevents.properties`
 - b) Change the `LicenseLogFile` parameter to point to the FLEXnet license log file. Even for Windows, you should use a slash rather than a backslash.

For example,

 - UNIX: `LicenseLogFile=/file_path/lmgrd.log`
 - Windows: `C:/file_path/lmgrd.log`
 4. Start the Analytics node host by restarting the loader controller (plc) service.

```
perfadmin stop plc
```

```
perfadmin start plc
```

Enable FLEXnet Manager data loaders in your cluster

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

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

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

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

- For `csh` or `tcsh`:

```
setenv DISPLAY hostname:0.0
```
- For `sh`, `ksh`, or `bash`:

```
DISPLAY=hostname:0.0
```

```
export DISPLAY
```

where `hostname` is your local host.

3. Add a data source for the FLEXnet Manager Reporting database to the Analytics node.

The FLEXnet Manager Reporting data source is named `FNMReportDB`.

 - In UNIX, run `ANALYTICS_TOP/bin/dbconfig.sh add FNMReportDB`
 - In Windows, run `ANALYTICS_TOP\bin\dbconfig add FNMReportDB`
4. Modify the FLEXnet Manager data loader configuration file to point to the FLEXnet Manager servers.
 - a) Edit the FLEXnet Manager data loader configuration file.
 - UNIX: `ANALYTICS_TOP/conf/dataloader/fnmloader.properties`
 - Windows: `ANALYTICS_TOP\conf\dataloader\fnmloader.properties`

- b) Change the `DataSource` parameter to specify the name of the data source for the FLEXnet Manager Reporting database.

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

For the FLEXnet Manager Reporting database,

```
DataSource=FNMRptDB
```

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

The element should now resemble the following:

```
<DataLoader Name="fml oader" ... Enable="true" ... />
```

6. Start the Analytics node host by restarting the loader controller (plc) service.

```
perfadmin stop plc
```

```
perfadmin start plc
```

Install and configure the Platform FLEXnet Manager scripts

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

Check that the following are installed and functional:

- FLEXnet Manager, version 11 or later
- Cognos PowerPlay 7.1 MR3 or later
- Oracle 9i or 10g database
- Jasper reports

Configure Platform's FLEXnet Manager (FNM) scripts to work with the FLEXnet Manager data loaders and your LSF cluster.

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

Edit `PERF_TOP/conf/plc/plc_license.xml`.

- b) Enable the FLEXnet Manager data loader.

```
<DataLoader Name="fml oader" ... Enable="true" ... />
```

- c) Restart the plc service for your changes to take effect.

2. Deploy the esub script.

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

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

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

Configure the Analytics node

Note the following:

- SLSF_CLUSTER_NAME is the name of the LSF cluster where the job is running
 - SSUBMIT is the job submission time as the number of non-leap seconds from 00:00:00 UTC, January 1, 1970.
- b) If you do not have an existing script, you need to copy and deploy the `esub.fnm` script, which is located in the PERF samples directory:
- UNIX: `$PERF_TOP/samples`
 - Windows: `%PERF_TOP%\samples`
3. Deploy the job starter script.

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

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

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

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

badmin reconfig

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

bqueues -l queue_name

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

bqueues -l normal

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

```
LM_PROJECT=$SLM_PROJECT, $LSB_JOBID, $SLSF_JOBINDEX  
export LM_PROJECT
```

- e) If you do not have an existing script, you need to copy and deploy the `jstarter.fnm` script, which is located in the PERF samples directory:
- UNIX: `$PERF_TOP/samples`
 - Windows: `%PERF_TOP%\samples`
4. Deploy the FNM data loader to a data collection host.

Verify and configure the Analytics server

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

What you need to do

To check the status of your cubes, datamarts, and data transformers, Analytics should be running correctly in your cluster for at least one day. This ensures that Analytics collects a sufficient amount of data.

1. Log into the Analytics server host.
2. Launch the Analytics Console.
 - UNIX: ***ANALYTICS_TOP/bin/runconsole.sh***
 - Windows: Start > Programs > Platform Analytics Server > 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 Purger in the navigation tree and compare the data purger settings with your cluster data retention policies.

To view the data purger settings, right-click each purger task and select View Purger Properties.

For more details on the data purger retention period for specific data types, refer [Optimize the Analytics server](#) on page 71 to and navigate to *Change the data retention period*.
5. Click Scheduled Tasks in the navigation tree and modify the times that the scheduled tasks are running, if necessary.
6. If you configured Analytics to collect FLEXnet license data, enable the FLEXnet Manager cubes in the build cubes scheduled tasks.
 - a) Click Scheduled Tasks in the navigation tree.
 - b) In the main window, right-click BuildCubes and select Edit Scheduled Task.
 - c) Move the following cubes from the Disabled column to the Enabled column:
 - Daily FLEX License
 - Daily Workload License
 - d) In the main window, right-click BuildCubesHourly and select Edit Scheduled Task.
 - e) Move the following cubes from the Disabled column to the Enabled column:
 - Hourly FLEX License
 - Hourly Workload License
7. Verify the schedule of the raw data purger scheduled task.
 - a) Click Scheduled Tasks in the navigation tree.
 - b) In the main window, right-click PurgerRawData and select Edit Scheduled Task.
 - c) If the Enable Scheduling checkbox is clear, select it to enable scheduling.
 - d) Change the schedule of this task to suit your data management policies.
8. Click Events in the navigation tree and verify that there are no ERROR or FATAL events.
9. Verify the email notification settings.

While in Events, click Action > Notification to open the Event Notification dialog.
10. After Analytics has been running for more than one day, click Cubes in the navigation tree and check the status of your cubes.
 - a) For any cubes with a "failed" status, right-click the cube and select Build Cube to manually build the cube.

- b) If the manual cube build still failed, click Events in the navigation tree and look for events related to your cube.

To find events specific to your cube, right-click the main window, select Filter Events, and enter the name of your cube in the Keyword field.

11. After Analytics has been running for more than one day, click Datamart ETL Flows in the navigation tree and check the status of your datamarts.

- a) For any datamarts with a "failed" status, right-click the datamart and select Build Datamart to manually build the datamart.
- b) If the manual cube build still failed, click Events in the navigation tree and look for events related to your datamart.

To find events specific to your datamart, right-click the main window, select Filter Events, and enter the name of your datamart in the Keyword field.

12. After Analytics has been running for more than one day, verify that there are no errors in the data transformer log files.

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

- UNIX: `ANALYTICS_TOP/log/datatransformer`
- Windows: `ANALYTICS_TOP\log\datatransformer`

After configuring the Analytics server, you can optionally optimize the Analytics node.

Verify and configure the Analytics server

Optimize the Analytics node

Modify the Analytics node to enhance performance.

What you need to do

The following steps are optional.

1. [Increase JVM memory](#) on page 66
2. [Disperse the Analytics node workload](#) on page 66
3. [Optimize specific data loaders](#) on page 69

After optimizing the Analytics node, you can optionally optimize the Analytics server.

Increase JVM memory

Increase the Java Virtual Machine (JVM) memory to at least 2 GB for the Analytics node instead of using the default value to prevent the JVM from running out of memory.

Note:

Java cannot use more than 1638 MB of memory on 32-bit platforms. If you are using 32-bit Windows, you can only increase the JVM memory to 1638 MB.

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

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

For example,

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

Disperse the Analytics node workload

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

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

Data loader category	Loader controller configuration file
Host-related	<ul style="list-style-type: none">• <code>plc_ego.xml</code>• <code>plc_coreutil.xml</code>
Job-related	<ul style="list-style-type: none">• <code>plc_lsf.xml</code>• <code>plc_bj_obs-sp012.xml</code>
Advanced job-related	<ul style="list-style-type: none">• <code>plc_lsf_advanced.xml</code>

Data loader category	Loader controller configuration file
License-related	<ul style="list-style-type: none"> pl c_l i cense. xml

- [Disperse the Analytics node workload to multiple hosts](#) on page 67
 Disperse the workload to multiple hosts if your Analytics node is experiencing performance issues and cannot handle all the workload by itself.
- [Disperse the Analytics node workload to multiple loader controllers in one node](#) on page 67
 Disperse the workload to multiple loader controllers within a host to reduce performance bottlenecks if your Analytics node host is powerful enough to handle the workload (for example, the host has at least four cores and at least 6 GB of memory).

Disperse the Analytics node workload to multiple hosts

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

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

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

Alternatively, you can have fewer nodes and still disperse workload to some degree. For example, you could have two nodes: your original node could handle host-related and job-related data loaders while another node could handle advanced job-related and license-related data loaders.

2. For each Analytics node, disable the data loading categories that the node will not be handling.
 - a) Navigate to the `ANALYTICS_TOP/conf/pl c` directory.
 - b) Move or rename the configuration files for the data loader categories that the Analytics node will no longer handle.

For example, if you want the Analytics node to only handle host-related data loaders, move or rename all configuration files except `pl c_ego. xml` and `pl c_coreuti l . xml`.

Disperse the Analytics node workload to multiple loader controllers in one node

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

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

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

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

For example, you can create `pl c_j ob` to handle job-related data loaders, `pl c_advanced_j ob` to handle advanced job-related data loaders, and `pl c_l i cense` to handle license-related data loaders, while the original loader controller (`pl c`) handles host-related data loaders.

Optimize the Analytics node

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

For each new loader controller, name the new file *wsm_PLC_NAME.conf*

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

- b) Edit each new *wsm_PLC_NAME.conf* file and specify the new loader controller name.

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

```
SERVICE_NAME=PLC_NAME
SERVICE_COMMAND=com.plattform.perf.dataloader.Main -f PLC_NAME.xml
LOG_PREFIX=PLC_NAME
```

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

```
SERVICE_NAME=plc_job
SERVICE_COMMAND=com.plattform.perf.dataloader.Main -f plc_job.xml
LOG_PREFIX=plc_job
```

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

- c) From the *ANALYTICS_TOP/conf* directory, create a new subdirectory for each new loader controller.

Name each directory *PLC_NAME*.

For example, create three new directories named *plc_job*, *plc_advanced_job*, and *plc_license*.

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

For example,

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

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

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

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

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

For example,

1. In *plc_job.xml*, change *Port* to from **4046** to **4047** and change *PLCDir* from **plc** to **plc_job**.
2. In *plc_advanced_job.xml*, change *Port* to from **4046** to **4048** and change *PLCDir* from **plc** to **plc_advanced_job**.

3. In `plc_license.xml`, change `Port` to from **4046** to **4049** and change `PLCDir` from `plc` to `plc_license`.
2. Restart the new loader controllers.

```
perfadmin stop all
```

```
perfadmin start all
```

Note:

To stop or start an individual data loader, use `perfadmin stop PLC_NAME` and `perfadmin start PLC_NAME`

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

Optimize specific data loaders

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

1. Optimize the FLEXnet usage data loader (`f1exliscusageloader`) to improve data loading time if it cannot finish loading the data within one sampling interval.

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

- a) Log into the Analytics node host that is running the FLEXnet usage data loader.
- b) Edit the `ANALYTICS_TOP/conf/dataloader/f1exliscusage.properties` file.
- c) Edit the `ThreadNumber` parameter to increase the number of threads for collecting data concurrently.

For example,

```
ThreadNumber=10
```

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

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

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

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

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

- a) Log into the Analytics node host that is running the LSF events data loader.
- b) Edit the `LSF_ENVDIR/lsbatch/cluster_name/configdir/lsb.params` file.
- c) In the `Parameters` section, define a new parameter named `MAX_EVENT_STREAM_SIZE`.

For example,

```
Begin Parameters
...
MAX_EVENT_STREAM_SIZE = 2048
...
```

Optimize the Analytics node

End Parameters

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

lsadmin reconfig

Optimize the Analytics server

Modify the Analytics server to enhance performance.

What you need to do

The following steps are optional.

1. [Modify the Cognos parameters](#) on page 72
2. [Modify the Cognos PowerPlay Enterprise Server parameters](#) on page 73
3. [Increase JVM memory](#) on page 74
4. [Split the cube build tasks to improve the cube build turnaround time](#) on page 75
5. [Improve the hourly build cubes task to show data from the last hour](#) on page 78
6. [Split the data transformer tasks to disperse workload](#) on page 80
7. [Optimize the data purger \(non-partitioned data schema only\)](#) on page 83
8. [Change the data retention period](#) on page 84

Modify the Cognos parameters

Edit the Cognos parameters to enhance the performance of the Cognos server and to improve cube building performance. The parameters you edit depend on the operating system of the Analytics server host.

- [Modify the Cognos parameters in a Windows host](#) on page 72
- [Modify the Cognos parameters in a Solaris host](#) on page 72

Modify the Cognos parameters in a Windows host

1. Log into the Analytics server host.
2. Navigate to the `cer5\bin` subdirectory of the Cognos installation directory.

If you installed Cognos to the default installation directory, navigate to `C:\Program Files\Cognos\cer5\bin`.

3. Edit the `cer5.ini` file and change the following parameter definitions:

```
WriteCacheSize=65536  
SortMemory=10240
```

4. Edit the `trnsfrmr.ini` file (or create a new file if it does not exist), and change (or add) the following parameter definitions:

```
[PowerPlay Transformer]  
MaxTransactionNum=800000  
MultiFileCubeThreshold=1000000
```

Modify the Cognos parameters in a Solaris host

1. Log into the Analytics server host.
2. Navigate to the `cer5/bin` subdirectory of the Cognos installation directory.

If you installed Cognos to the default installation directory, navigate to `/usr/cognos/cer5/bin`.

3. Edit the `cer5.ini` file and change the following parameter definitions:

```
SORT_MEMORY=10240
```

4. Edit the `trnsfrmr.rc` file (or create a new file if it does not exist), and change (or add) the following parameter definitions:

```
[PowerPlay Transformer]
```



```
MaxTransactionNum=800000
MultiFileCubeThreshold=1000000
```

- Optional. Edit `ANALYTICS_TOP/bin/cr_cube.sh` and change the `PPDS_WRITE_MEMORY` environment variable definition to another value.

The default parameter definition (65536) is already at the optimal value, so you should not need to change this value.

Modify the Cognos PowerPlay Enterprise Server parameters

Edit the Cognos PowerPlay Enterprise Server parameters to improve cube access performance.

- Log into the Analytics server host.
- If you connected to the intended Analytics server host via `telnet` and are running `xserver` on a local host, set your display environment.

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

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

- For `csch` or `tcsh`:

```
setenv DISPLAY hostname:0.0
```

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

```
DISPLAY=hostname:0.0
```

```
export DISPLAY
```

- Start the Cognos PowerPlay Enterprise Server (PPES).
 - UNIX: Navigate to the `cer5/bin` subdirectory of the Cognos installation directory and run `./ppsrvadm.sh`.
 - Windows (Cognos 7.4 MR3): Navigate to the `cer5\bin` subdirectory of the Cognos installation directory and run `PPSrvAdmin.exe`.
 - Windows (Cognos 7.4 MR4): Select Start > Programs > IBM Cognos Series 7 Version *number* > IBM Cognos Server Administration > OLAP Cubes and Reports.
- Specify the user ID and password to connect to the Cognos PPES.
- Right-click the host name from the left navigation tree and select Properties.
- From the Server Properties dialog, select Cube Settings.
- Change the definition for the following parameters, then click Apply:

Parameter name	Value
Queued request timeout(s)	3000
Maximum processes	100
Request timeout(s)	9000
Recycle Time(min)	14400

- From the Server Properties dialog, select Report Settings.
- Change the definition for the following parameters, then click Apply:

Parameter name	Value
Queued request timeout(s)	3000
Maximum processes	100
Recycle Time(min)	14400
Report timeout(min)	500

10. Restart the Cognos PPES to apply the changes.

Increase JVM memory

Increase the Java Virtual Machine (JVM) memory for the data transformer and cube and datamart building. The file and parameters that you modify depend on the operating system of the Analytics server host.

- [Increase JVM memory in a Windows host](#) on page 74
- [Increase JVM memory in a Solaris host](#) on page 75

Increase JVM memory in a Windows host

Note:

Java cannot use more than 1638 MB of memory on 32-bit platforms. If you are using 32-bit Windows, you can only increase the JVM memory to 1638 MB.

1. Log into the Analytics server host.
2. Increase the JVM memory for the data transformer.

Edit the `ANALYTI_CS_TOP\bin\runETL.bat` file and change the Java starting command line to the following:

- Windows 32-bit (you can only change this to a maximum of 1638 MB):
`"%JAVA_HOME%\bin\java.exe" -Xms64m -Xmx1638m`
- Windows 64-bit (you can change this to 2 GB or more):
`"%JAVA_HOME%\bin\java.exe" -Xms64m -Xmx2048m`

Note:

If you have a large cluster (with at least 1000 hosts), you should increase the JVM memory to at least 4 GB. Therefore, you must be using a Windows 64-bit host, and you should change the Java starting command to use at least 4096 MB:

```
"%JAVA_HOME%\bin\java.exe" -Xms64m -Xmx4096m
```

3. Increase the JVM memory for cube and datamart building.

Edit the `ANALYTI_CS_TOP\bin\rundatamart.bat` and change the Java starting command line to the following:

- Windows 32-bit (you can only change this to a maximum of 1638 MB):
`"%JAVA_HOME%\bin\java.exe" -Xms64m -Xmx1638m`

- Windows 64-bit (you can change this to 2 GB or more):
`"%JAVA_HOME%\bin\java.exe" -Xms64m -Xmx2048m`

Increase JVM memory in a Solaris host

1. Log into the Analytics server host.
2. Increase the JVM memory for the data transformer.

Edit the `ANALYTICS_TOP/bin/ex_runETL.sh` file and change the Java starting command line to use 2 GB or more memory, as follows:

```
$JAVA_HOME/bin/java -Xms64m -Xmx2048m
```

Note:

If you have a large cluster (with at least 1000 hosts), you should increase the JVM memory to at least 4 GB. Therefore, you should change the Java starting command to use at least 4096 MB:

```
$JAVA_HOME/bin/java -Xms64m -Xmx4096m
```

3. Increase the JVM memory for the data transformer.

Edit the `ANALYTICS_TOP/bin/rundatamart.sh` file and change the Java starting command line to use 2 GB or more memory, as follows:

```
$JAVA_HOME/bin/java -Xms64m -Xmx2048m
```

Split the cube build tasks to improve the cube build turnaround time

By default, there are two scheduled tasks for building OLAP cubes that are scheduled to run every day. The cubes are published to Cognos Upfront after they are successfully built. In large cluster at maximum capacity (for example, 5 000 hosts with 3 000 000 finished jobs per day) the cube build tasks may take over 10 hours to complete. This means that even after running overnight, you might not see the updated cubes in the morning.

Before performing this optimization, you must determine how many cubes your Analytics server can run in parallel, which you can calculate by using the number of cores available in your host:

$(\text{Number of parallel build cubes}) = (\text{Number of cores available}) - 4$

The four cores that you subtract from the calculation are due to the four cores that are assigned to other Analytics tasks. Therefore, you cannot build two cubes in parallel unless the Analytics server has six or more cores available. This determines the maximum number of hourly cube build tasks and daily cube build tasks that you should add.

For example, if your Analytics server has eight cores available, you can build four cubes in parallel, and therefore you should not add more than four hourly cube build tasks and four daily cube build tasks.

Tip:

To improve the turnaround time of the hourly build cube tasks to an hour so that the hourly build cube tasks show data from the last hour, you may need to increase the number of parallel build cube tasks and therefore you may require more cores in the Analytics server. For an example of the number of cores you need in your Analytics server to be able to obtain

a one hour turnaround time for hourly build cube tasks, refer to [Improve the hourly build cubes task to show data from the last hour](#) on page 78. However, you should still follow the steps described here in order to split the cube build tasks for both hourly and daily cubes.

To improve the cube build turnaround time and see the updated cubes in the morning, split the two cube build tasks to at least four (two hourly cube build tasks and two daily cube build tasks). The following is a recommended format for splitting your cubes into two hourly cube build tasks and two daily cube build tasks:

Task name	Cube name
BuildCubesHourly1	Hourly Host Availability
	Hourly Workload
	Hourly Job Throughput
	Cluster Capacity
BuildCubesHourly2	Health Check
	Hourly Hardware
	Hourly Job Statistics
	Jobs vs Slot Utilization
	Hourly Job Resource Consumption
	Hourly Jobs by Pending Reason
	Hourly Shared Resource
	Hardware Histogram
	Hourly License Usage by User
	Hourly License Utilization
	Hourly License Consumption
	Hourly License Utilization by Feature
	Hourly License Denial
	Hourly License Denial by Feature
License Usage by Server Histogram	
License Usage by Feature Histogram	
BuildCubesDaily1	Daily Workload

Task name	Cube name
BuildCubesDaily2	Daily Hardware
	Daily Shared Resource
	Daily Job Throughput
	Daily Job Statistics
	Daily Job Resource Consumption
	Daily Jobs by Pending Reason
	Daily License Utilization
	Daily License Consumption
	Daily License Usage by User
	Daily License Denial
	Daily License Utilization by Feature
	Denial License Denial by Feature
	Daily Host Availability

1. Log into the Analytics server host.
2. Launch the Analytics Console.
 - UNIX: **ANALYTICS_TOP/bin/runconsole.sh**
 - Windows: Start > Programs > Platform Analytics Server > Platform Analytics Console
3. For each cube that the Analytics server will build in parallel, create a new scheduled task to run every hour starting from 2:00 a.m. for hourly cubes or to run every day at 3:30 a.m. for daily cubes.

For example, repeat the following steps to create two hourly scheduled tasks named `BuildCubesHourly1` and `BuildCubesHourly2`, and two daily scheduled tasks named `BuildCubes1` and `BuildCubes2`.

 - a) Click Scheduled Tasks in the navigation tree.
 - b) Right-click the main window and select Add Scheduled Task.
 - c) Complete the required fields for the new task.
 - Scheduled Task: Specify the name of this task.
 - Script File: Specify **bin/buildcubes.js**.
 - Script Function: Specify **doit**.
 - d) Enable the new scheduled task.
 - e) Change the new scheduled task to start at 2:00 a.m. for hourly cubes or 3:30 a.m. for daily cubes.
 - f) Select the Run every: field and specify 1 hour for hourly cubes or 1 day for daily cubes.
 - g) In the Cubes panel, select the cubes you want to build with this task and move them from Disabled to Enabled.
 - h) To save your changes and close the window, click OK.
4. From the `BuildCubesHourly` scheduled task, disable the hourly cubes that the new scheduled tasks will build.

- a) Click **Scheduled Tasks** in the navigation tree.
- b) Right-click the **Buil dCubesHourly** scheduled task in the main window and select **Edit Scheduled Task**.

The **Scheduled Task** window for **Buil dCubesHourly** displays.

- c) In the **Cubes** panel, select the hourly cubes that the new scheduled tasks will build move them from **Enabled** to **Disabled**.
5. From the **Buil dCubes** scheduled task, disable the daily cubes that the new scheduled tasks will build.
 - a) Click **Scheduled Tasks** in the navigation tree.
 - b) Right-click the **Buil dCubes** scheduled task in the main window and select **Edit Scheduled Task**.

The **Scheduled Task** window for **Buil dCubes** displays.

- c) In the **Cubes** panel, select the daily cubes that the new scheduled tasks will build move them from **Enabled** to **Disabled**.
6. Exit from the **Analytics Console**.
 7. Edit the **pi . conf** file.

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

8. Increase the specified value for the **BatchBuil dThreshold** parameter to the number of cubes that the Analytics server will build in parallel.

This parameter, which can have a value from 1 to 6, specifies the number of cubes that can be built in parallel.

For example, to build four cubes in parallel,

BatchBuildThreshold = 4

Improve the hourly build cubes task to show data from the last hour

By default, the scheduled task for building hourly OLAP cubes is scheduled to run every day. The cubes are published to Cognos Upfront after they are successfully built. Since the scheduled task runs every day, the hourly cubes only show the data from the last day. If you want the hourly cubes to always show data from the last hour, you should separate hourly cubes from the build cubes task into new scheduled tasks, and configure the Analytics server to run these tasks in parallel.

Before performing this optimization, you must determine how many cubes your Analytics server can run in parallel, which you can calculate by using the number of cores available in your host:

(Number of parallel build cubes) = (Number of cores available) - 4

The four cores that you subtract from the calculation are due to the four cores that are assigned to other Analytics tasks. Therefore, you cannot build two cubes in parallel unless the Analytics server has six or more cores available. To obtain a one hour turnaround time for all hourly cubes, the recommended value for this parameter will vary.

Tip:

To improve the cube build turnaround time for both the hourly and daily cube build tasks, refer to [Split the cube build tasks to improve the cube build turnaround time](#) on page 75. However, to obtain a one hour

turnaround time for hourly cube build tasks, your Analytics server must have at least the number of cores described here.

For example, in a cluster with 2 000 hosts and 1 000 000 finished jobs per day, you need to build at least four cubes in parallel, which means you need to have at least eight cores available. In a cluster with 5 000 hosts and 3 000 000 finished jobs per day, you need to build at least five cubes in parallel, which means you need to have at least nine cores available.

1. Log into the Analytics server host.
2. Launch the Analytics Console.
 - UNIX: **`ANALYTICS_TOP/bin/runconsole.sh`**
 - Windows: Start > Programs > Platform Analytics Server > Platform Analytics Console
3. For each cube that the Analytics server will build in parallel, create a new scheduled task to run every hour starting from 2:00 a.m.

For example, to build four cubes in parallel, repeat the following steps to create four scheduled tasks named `Bui 1 dCubesHour1 y1`, `Bui 1 dCubesHour1 y2`, `Bui 1 dCubesHour1 y3`, and `Bui 1 dCubesHour1 y4`.

- a) Click Scheduled Tasks in the navigation tree.
- b) Right-click the main window and select Add Scheduled Task.
- c) Complete the required fields for the new task.
 - Scheduled Task: Specify the name of this task.
 - Script File: Specify **`bin/buildcubes.js`**.
 - Script Function: Specify **`doit`**.
- d) Enable the new scheduled task.
- e) Change the new scheduled task to start at 2:00 a.m.
- f) Select the Run every: field and specify 1 hour.
- g) In the Cubes panel, select the hourly cubes you want to build with this task and move them from Disabled to Enabled.

Tip:

Distribute the hourly cubes evenly among all the new scheduled tasks.

- h) To save your changes and close the window, click OK.
4. From the `Bui 1 dCubesHour1 y` scheduled task, disable the hourly cubes that the new scheduled tasks will build.
 - a) Click Scheduled Tasks in the navigation tree.
 - b) Right-click the `Bui 1 dCubesHour1 y` scheduled task in the main window and select Edit Scheduled Task.

The Scheduled Task window for `Bui 1 dCubesHour1 y` displays.
 - c) In the Cubes panel, select the hourly cubes that the new scheduled tasks will build move them from Enabled to Disabled.
5. Exit from the Analytics Console.
6. Edit the `pi . conf` file.
 - UNIX: **`ANALYTICS_TOP/conf/pi . conf`**
 - Windows: **`ANALYTICS_TOP\conf\pi . conf`**

- Increase the specified value for the `BatchBuildThreshold` parameter to the number of cubes that the Analytics server will build in parallel.

This parameter, which can have a value from 1 to 6, specifies the number of cubes that can be built in parallel.

For example, to build four cubes in parallel,

BatchBuildThreshold = 4

Split the data transformer tasks to disperse workload

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

The following is a recommended format for splitting your data transformers into seven tasks. The examples will make use of this table (specifically, with Task 1). In the following example, Tasks 1 to 3 are hourly scheduled tasks, while Tasks 4 and 5 are daily scheduled tasks.

Task	Data transformer name	Data flow entry
1	IsfJobResUsageByGroupHourly	datatransformer/flow/bjobs-sp012/hourly/jobdetailresusage/main_job_resusage_bygroup.xml
	IsfJobSlotsUsageByGroupHourly	datatransformer/flow/bjobs-sp012/hourly/jobdetailresusage/main_job_slotsusage_bygroup.xml
	IsfJobStatisticsHourly	datatransformer/flow/bjobs-sp012/hourly/jobdetailresusage/main_job_slotsusage_bygroup.xml
2	IsfJobmartHourly	datatransformer/flow/jobrelateddataaggregation/hourly/jobmart/main_jobmart_etl.xml
	IsfJobThroughputHourly	datatransformer/flow/jobrelateddataaggregation/hourly/jobthroughput/main_job_throughput_etl.xml

Task	Data transformer name	Data flow entry
3	hostPropertiesHourly	datatransformer/flow/hostproperties/hourly/ main_host_properties.xml
	hostMetricsHourly	datatransformer/flow/hostmetrics/hourly/ main_host_metrics_hourly.xml
	hostGroupConfHourly	datatransformer/flow/lsfhostgroupconf/hourly/ main_host_group_conf_hourly.xml
	userGroupConfHourly	datatransformer/flow/lsfusergroupconf/hourly/ main_userGroupConfHourly.xml
	lsfSharedResourceUsageHourly	datatransformer/flow/lsfsharedresourceusage/hourly/ main_sharedResourceUsageHourly.xml
	lsfSharedResourcePropertiesHourly	datatransformer/flow/lsfsharedresourceproperties/hourly/ main_sharedResourceProperties.xml
	lsfHostStatusHourly	datatransformer/flow/lsfhoststatus/hourly/ main_lsf_host_status_hourly.xml
	lsfJobExecProcHourly	datatransformer/flow/jobrelateddataaggregation/hourly/ jobexecproc/main_job_execproc_etl.xml
	lsfJobSlotUsageByHost	datatransformer/flow/jobrelateddataaggregation/hourly/ jobslotusagebyhost/main_jobslotusagebyhost_etl.xml
	lsfJobSlotUsageByQueue	datatransformer/flow/jobrelateddataaggregation/hourly/ jobslotusagebyqueue/main_jobslotusagebyqueue_etl.xml
	lsfJobPendingReasonHourly	datatransformer/flow/lsfpendingreason/hourly/ main_job_pending_reason_hourly.xml
	lsfDetailPendingReasonHourly	datatransformer/flow/lsfdetailpendreason/hourly/hostreason/ main_detail_pending_reason_hourly.xml
	cluster_slots	datatransformer/flow/jobrelateddataaggregation/hourly/ job_v_slotutil/main_cluster_slots_etl.xml
	jobs_v_slotutil	datatransformer/flow/jobrelateddataaggregation/hourly/ job_v_slotutil/main_job_v_slotutil_etl.xml
	hostMetricsHistogram	datatransformer/flow/hardware_and_license_histogram-sp011/ hostmetrics_histogram/main_host_metrics_hourly.xml
hardwareCoreUtilHourly	datatransformer/flow/hardware_with_core_util-sp006/hourly/ main_hardware_core_util_hourly.xml	
4	lsfJobmartDaily	datatransformer/flow/jobrelateddataaggregation/daily/jobmart/ main_jobmart_etl_daily.xml
	lsfJobThroughputDaily	datatransformer/flow/jobrelateddataaggregation/daily/ jobthroughput/main_jobthroughput_etl_daily.xml
	lsfJobmartHPCDaily	datatransformer/flow/jobrelateddataaggregation/daily/jobmart/ main_hpc_jobmart_etl_daily.xml

Task	Data transformer name	Data flow entry
5	hostMetricsDaily	datatransformer/flow/hostmetrics/daily/main_host_metrics_daily.xml
	hostNotBuiltinMetricsDaily	datatransformer/flow/hostmetrics/daily/main_host_notbuiltin_metrics_daily.xml
	hostGroupConfDaily	datatransformer/flow/lsfhostgroupconf/daily/main_host_group_conf_daily.xml
	IsfHostStatusDaily	datatransformer/flow/lsfhoststatus/daily/main_lsf_host_status_daily.xml
	IsfSharedResourcesDaily	datatransformer/flow/lsfsharedresourceusage/daily/main_sharedResourceUsageDaily.xml
	IsfJobResUsageCostDaily	datatransformer/flow/jobrelateddataaggregation/daily/jobresusagecost/main_jobresusagecost_etl_daily.xml
	IsfJobStatisticsDaily	datatransformer/flow/jobrelateddataaggregation/daily/jobstatistics/main_jobnumberofjobs_etl_daily.xml
	IsfJobResUsageByGroupDaily	datatransformer/flow/jobrelateddataaggregation/daily/jobresusagebygroup/main_jobresusagebygroup_etl_daily.xml
	IsfJobSlotUsageByGroupDaily	datatransformer/flow/jobrelateddataaggregation/daily/jobslotsusagebygroup/main_jobslotusagebygroup_etl_daily.xml
	IsfJobPendingReasonDaily	datatransformer/flow/lsfpendingreason/daily/main_job_pending_reason_daily.xml
	IsfDetailPendingReasonJobNum Daily	datatransformer/flow/lsfdetailpendreason/daily/hostreason/jobnumber/main_detail_pend_reason_jobnum_daily.xml
	IsfDetailPendingReasonTimeDaily	datatransformer/flow/lsfdetailpendreason/daily/hostreason/pendtime/main_detail_pend_reason_pendtime_daily.xml
	hardwareCoreUtilDaily	datatransformer/flow/hardware_with_core_util-sp006/daily/main_hardware_core_util_daily.xml

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

For example, if you are creating Task 1 from the table with the recommended format of splitting data transformers, specify the following:

- Scheduled Task: Specify Task1 as the name of the scheduled task.
 - Script File: Specify `bin/dataagghourly.js` as the path to the script file.
 - Script Function: Specify `doit` as the script function.
- e) Enable the new scheduled task that you created.
3. In the `tasks` subdirectory of `ANALYTICS_TOP`, create a new directory with the same name as the name of the new scheduled task and navigate to the new directory.

For example, for Task 1 on a UNIX host,

```
cd ANALYTICS_TOP/tasks
```

```
mkdir Task1
```

```
cd Task1
```

4. From the new directory, create a text file of any name with the `.tsk` extension.

For example, create `task1.tsk`.

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

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

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

```
# IsfJobResUsageByGroupHourly
datatransformer/flow/bjobs-sp012/hourly/jobdetailresusage/main_job_resusage_bygroup.xml
# IsfJobSlotsUsageByGroupHourly
datatransformer/flow/bjobs-sp012/hourly/jobdetailresusage/main_job_slotsusage_bygroup.xml
# IsfJobStatisticsHourly
datatransformer/flow/bjobs-sp012/hourly/jobstatistics/main_jobstatistics_etl.xml
```

Optimize the data purger (non-partitioned data schema only)

To optimize the data purger to the recommended level, you must be using a cluster with less than 1000 hosts and have approximately 40 GB of UNDO tablespace in the database.

If you are using a non-partitioned data schema in a cluster with less than 1000 hosts, you can change the data purger batch size to improve the performance of the data purger. If you have a cluster of more than 1000 hosts, you should use partitioned data schema.

If you are using a partitioned data schema, your data purger is already operating at optimal levels, and the following changes will not improve its performance.

1. Log into the Analytics server host.
2. Edit the `pi.conf` file.
 - UNIX: `ANALYTICS_TOP/conf/pi.conf`
 - Windows: `ANALYTICS_TOP\conf\pi.conf`
3. Change the following parameter definition to the recommended level:

```
PURGER_BATCH_SIZE=1000000
```

To set the parameter to this level, you must be using a cluster with less than 1000 hosts and have approximately 40 GB of UNDO tablespace in the database.

If your cluster is greater than 1000 hosts, you should use a partitioned data schema, as the data purger cannot function at optimal performance with non-partitioned data schema for a cluster of this size.

Change the data retention period

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

The steps you take to change the data retention period depend on whether you are using partitioning for the data schema.

- [Change the data retention period for non-partitioned data schema](#) on page 84
- [Change the data retention period for partitioned data schema](#) on page 85

Change the data retention period for non-partitioned data schema

If you are using non-partitioned data schema, the data purger consists of two scheduled tasks (PurgeRawData and PurgeAggData), which are disabled by default. To purge out-of-date data, enable these two scheduled tasks using the Analytics Console.

The following table describes the default data retention period grouped by data type. Each configuration file is located in the purger configuration directory:

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

Data type	Default data retention period (days)	Configuration files
Raw data	14	<ul style="list-style-type: none"> • bjobs-sp012.xml • purger_hardware_with_core_util-sp006_rawdata.xml • purger_license_rawdata.xml • purger_lsf_advanced_rawdata.xml • purger_lsf_basic_rawdata.xml
Job- and license-related hourly data	84	<ul style="list-style-type: none"> • purger_hardware_with_core_util-sp006_hourlydata.xml • purger_histogram-sp011.xml • purger_license_hourlydata.xml • purger_lsf_job_hourlydata.xml
Host-related hourly data	14	<ul style="list-style-type: none"> • purger_lsf_host_hourlydata.xml
Daily data	182	<ul style="list-style-type: none"> • purger_hardware_with_core_util-sp006_dailydata.xml • purger_license_dailydata.xml • purger_lsf_dailydata.xml

1. Log into the Analytics server host.
2. Enable the PurgeRawData and PurgeAggData scheduled tasks using the Analytics Console.

- a) Launch the Analytics Console.
 - UNIX: ***ANALYTICS_TOP/bin/runconsole.sh***
 - Windows: Start > Programs > Platform Analytics Server > Platform Analytics Console
 - b) Click Scheduled Tasks in the navigation tree.
 - c) Enable the PurgeRawData and PurgeAggData scheduled tasks.
3. Edit each configuration file for the data type that you want to change.

Within each configuration file, navigate to the TableList element and change the Duration attribute to the duration that you want.

For example, to change the data retention period to 28 days, change the Duration attribute to the following,

```
<TableList ... Duration="28" ... />
```

Change the data retention period for partitioned data schema

If you are using partitioned data schema, the data purger consists of two scheduled tasks (RawTablePartition and WiTablePartition), which are enabled by default.

The following table describes the default data retention period grouped by data type. Each configuration file is located in the bin directory:

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

Data type	Default data retention period (days)	Configuration file	Command to purge the data
Host-related raw data	5	autopartition_raw.js	PA_PARTITION.DROP_PARTITION('RAW_HOST', 1, SYSDATE + 2, 7, 1)
Job- and license-related raw data	14	autopartition_raw.js	PA_PARTITION.DROP_PARTITION('RAW', 1, SYSDATE + 2, 16, 1)
Host-related hourly data	18	autopartition_wi.js	PA_PARTITION.DROP_PARTITION('HOURLY_HOST', 2, SYSDATE + 2, 10, 2)
Job- and license-related hourly data	98	autopartition_wi.js	PA_PARTITION.DROP_PARTITION('HOURLY', 2, SYSDATE + 2, 14, 7)
Daily data	222	autopartition_wi.js	PA_PARTITION.DROP_PARTITION('DAILY', 2, SYSDATE + 2, 8, 28)

In the commands to purge the data, the fourth parameter specifies the number of partitions for each table to be kept, and the fifth parameter specifies the number of days of data that is kept; the data purger deletes all remaining older data. Therefore, for host-related raw data, the default is 16 partitions kept and the data retention period is one day.

To change the default retention period for each data type, edit the configuration file and navigate to its corresponding command to purge the data. Change the fifth parameter to the new retention period.

Optimize the Analytics server

Remove Analytics software from the host

Remove Analytics from a UNIX or Windows host.

What you need to do

- [Remove Analytics software from a UNIX host](#) on page 88
- [Remove Analytics software from a Windows host](#) on page 88

Remove Analytics software from a UNIX host

1. Log into the Analytics host with the same account that you used to install Analytics.
2. Navigate to the *ANALYTICS_TOP*/uninstall directory.
3. Run the uninstall package.

`./uninstaller.bin`

To remove Analytics cubes from Cognos, refer to your Cognos documentation.

Remove Analytics software from a Windows host

1. Log into the Analytics host with the same account that you used to install Analytics.
2. Run the Analytics uninstaller.
 - a) Select Start > Settings > Control Panel > Add/Remove Programs.
 - b) Select your Platform Analytics installation and click Remove.
 - c) If you receive a message that the program has already been uninstalled, manually run the Analytics uninstaller.
 1. Navigate to the *ANALYTICS_TOP*\uninstall directory.
 2. Run the uninstall package.

`uninstaller.exe`

Restart your host to complete the uninstallation process.

To remove Analytics cubes from Cognos, refer to your Cognos documentation.

Index

A

- Adobe Acrobat Reader 26, 27
- Analytics Console 62
- ANALYTICS_TOP 49, 53
- Apache web server
 - accessing 32, 39
 - configuring 31, 38
- architecture 6

C

- Cognos 25
 - Access Manager
 - configuring 30, 37
 - Configuration Manager
 - Solaris 28, 36
 - Windows 33, 40
 - directory server
 - installing 25, 26
 - instance 28, 33
 - PowerPlay Enterprise Server
 - configuring 30, 37
 - installing 26, 27
 - securing 32, 39
 - testing 43
 - PowerPlay Transformer Edition
 - installing 26, 27
 - server groups
 - configuring 34, 40
 - Upfront
 - configuring 30, 37
 - testing 43
 - Upfront File Manager
 - configuring 31, 38

- cubes 62
 - default publishing mode 42

D

- data loaders
 - log files 50
 - verifying 50

- data purger 62
- data sources
 - FNMReportDB 58
- data transformers 63
- database
 - ports 5
 - schema 19
- datamarts 63
- directory server
 - configuring 29, 33, 36, 40
 - installing 25, 26
 - instance 28, 33
- display environment
 - setting 28, 36, 48, 52, 58, 73

E

- environment
 - display
 - setting 28, 36, 48, 52, 58, 73
 - LSF
 - sourcing 49, 56, 57
 - PERF
 - sourcing 49
- events
 - Analytics Console 62
 - notification 53, 62

F

- files
 - installation
 - Analytics 14
 - JDBC driver 14
 - license 9, 14
- FLEXnet
 - data loaders
 - configuring 57, 58
 - enabling 57, 58

- data source 58
- FNMRReportDB 58

I

- IIS web server
 - accessing 35, 42
 - configuring 34, 41

J

- JDBC driver 14, 48, 52

L

- license 9
 - FLEXnet
 - data loaders 57, 58
- loader controller
 - ports 49
- LSF_ENVDIR 49

N

- node 3
 - Analytics Console 62
 - ports 4
 - uninstalling 88

O

- Oracle
 - database
 - ports 5
 - schema 19
 - JDBC driver 48, 52

P

- PERF
 - services
 - starting 50
 - stopping 56
 - verifying 49
- ports 4, 24

- database 5
 - loader controller 49
 - node 4
 - server 4, 53

- purger 62

S

- scheduled tasks 62

- server 3
 - configuring 62
 - ports 4, 53
 - uninstalling 88
 - verifying 62

- services

- starting 50
- stopping 56
- verifying 49

- source environment

- LSF 49, 56, 57
- PERF 49

- Sun Java

- Directory Server

- configuring 29, 33, 36, 40

- directory server instance 28, 33

- installing 25, 26

- System Server Console

- Solaris 28

- Windows 33

- system architecture 6

U

- uninstall 88

W

- web server

- Apache (Solaris)

- accessing 32, 39

- configuring 31, 38

- IIS (Windows)

- accessing 35, 42

- configuring 34, 41