

Platform RTM
Version 9.1

Administering



Platform RTM
Version 9.1

Administering



Note

Before using this information and the product it supports, read the information in “Notices” on page 79.

First edition

This edition applies to version 9, release 1 of Platform RTM (product number 5725-G82) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1992, 2013.**

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

Contents

Figures	v
--------------------------	----------

Tables	vii
-------------------------	------------

IBM Platform RTM and Cacti	ix
---	-----------

Chapter 1. About Platform RTM.	1
---	----------

IBM Platform RTM and Cacti.	2
-------------------------------------	---

Chapter 2. Monitoring clusters with Platform RTM	3
---	----------

Configure cluster management	3
Management section.	3
Grid Management section	6
Configure cluster interaction	8
Grid Settings page	9
Cacti Settings page	12
Control an LSF cluster.	13
Run grid control commands on an LSF cluster.	14
Enforce cluster action control comments	14
Set mandatory cluster action control comments	14
Comment on cluster action control by host, queue, or cluster.	14
Configure idle job detection for cluster	15
Configure idle job detection (Grid settings).	15
Configure memory violation detection	16
Add or edit LSF clusters for Platform RTM to monitor.	17
Add or edit clusters	17
Add the RTM host to the LSF cluster as an LSF client	19
Add all hosts in the LSF cluster to RTM as devices	20
View LSF cluster and job information.	20
Dashboards section.	21
Job Info section	23
User/Group Info section	33
Load Info section	34
Host Info section	35
Reports section	37
Settings tab (from the Grid page)	38
View time series graphs for clusters and hosts.	39
Viewing modes	39
Settings tab (from the Graph page)	40
View non-time series charts for clusters and hosts	40

Chapter 3. Working with application license servers	43
--	-----------

Relationship between Platform RTM and license servers	43
License Management section.	43
Configure license pollers	44
Set up remote license pollers	44

Add or edit license pollers to the Platform RTM Console.	44
Add or edit license servers for Platform RTM to monitor.	45
Monitor license servers	46
View the status of the license services	46
View reports on license servers and projects	47
Create graphs to monitor your license server	48
View time series graphs for license servers	48
Viewing non-time series charts for license servers	49

Chapter 4. Setting up your email notifications for alerts	51
--	-----------

Working with thresholds and alerts	51
Configure thresholds and alerts.	51
View threshold and alert information.	53
Acknowledge triggered alerts for thresholds	54
Working with grid alerts and templates	55
Creating an Alert Template	56
Creating an Alert	57
Viewing Alert Log	58
Syslogs alerts.	58

Chapter 5. External Load Indices manager (ELIM) Graphs	59
---	-----------

Creating an ELIM Template	59
Creating an ELIM Instance	59
Creating a Graph for an ELIM Instance	59
GPFS ELIM	60
Configuring ELIM Script	60
Configuring LSF cluster	60
Creating graphs for GPFS ELIM	61

Chapter 6. Managing the metadata plug-in	63
---	-----------

Define metadata XML configuration file	63
Define metadata using CSV file.	65
Enabling metadata settings in Realm Permissions.	65
Importing metadata XML or CSV configuration file via RTM Console	65
Viewing metadata details.	66

Chapter 7. Performance and maintenance	67
---	-----------

Performance tuning for large clusters.	67
Store rrd files on a separate disk	67
Increase the database memory	67
Enable on-demand rrd file updating for systems with heavy disk I/O	67
Configure concurrent poller processes	68
Enable database record partitioning	68
Enable Ajax filter	70
Maintaining Database	70
Set up record partitioning in the database	70

Back up and restore the Platform RTM system settings	71
Setting up a remote database host	77

Notices	79
Trademarks	81

Figures

1. Overview of Platform RTM 1

Tables

- | | | | | | |
|----|---|---|----|------------------------------|----|
| 1. | Names, tags, and descriptions of available placeholders | 5 | 2. | LIM status icons | 22 |
| | | | 3. | Batch status icons | 22 |

IBM Platform RTM and Cacti

IBM® Platform RTM caters to three user groups who are each responsible for monitoring and reporting in the LSF® environment. LSF administrators, who are responsible for monitoring and maintaining the LSF clusters and application license servers, are the most common users of Platform RTM.

This documentation assumes that you are familiar with Cacti. For an introduction to Cacti itself, and for information specific to Cacti, refer to the Cacti documentation at cacti.net/documentation.php.

Chapter 1. About Platform RTM

IBM Platform RTM is an operational dashboard for IBM Platform LSF environments that provides comprehensive workload monitoring, reporting, and management. It makes cluster administrators more efficient in their day-to-day activities and provides the information and tools that are needed to improve cluster efficiency, enable better user productivity, and contain or reduce costs.

Unlike other monitoring tools that focus on just one facet of cluster monitoring, Platform RTM provides a complete, integrated monitoring facility that is designed specifically for Platform LSF environments. Multiple clusters can be monitored easily and effectively through a single intuitive interface.

Figure 1 shows a high-level view of Platform RTM in an LSF environment.

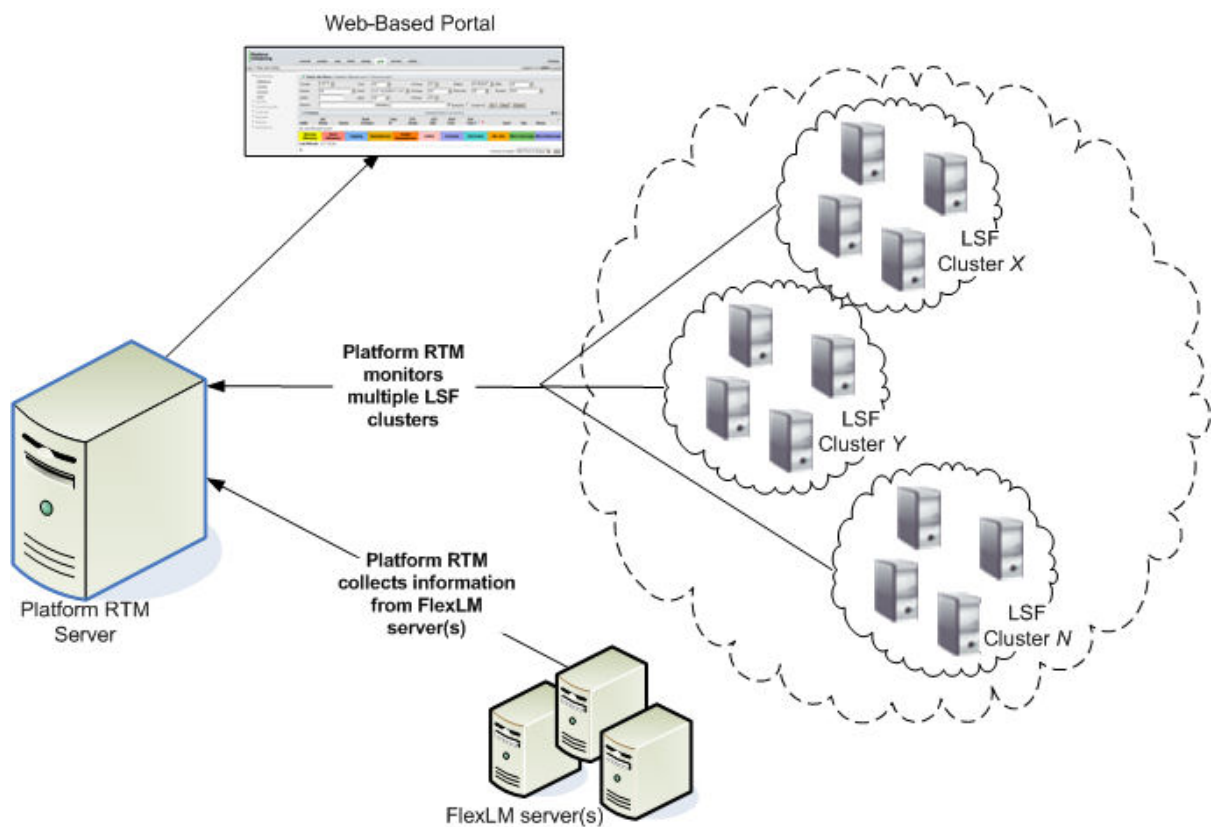


Figure 1. Overview of Platform RTM

Platform RTM is used to monitor and graph LSF resources (including networks, disks, applications, and others) in a cluster. In graph or table formats, Platform RTM displays resource-related information such as the number of jobs that are submitted, the details of individual jobs (like load average, cpu usage, job owner), or the hosts on which the jobs ran.

FLEXlm is a third-party license manager that is monitored by Platform RTM. FLEXlm allows licenses to remain on the network instead of a specific host. Any registered host in the cluster can then access and use a specific licensed application, increasing usage efficiency.

IBM Platform RTM and Cacti

Cacti is a complete RRDTool-based graphing solution that is developed by The Cacti Group. IBM Platform RTM uses Cacti as a rich graphical user interface framework to provide monitoring, reporting, and alerting functions specifically for the LSF environment. The LSF capabilities are included as a Cacti plug-in so that when used together, IBM Platform RTM can offer LSF-specific monitoring and reporting capabilities. These features are in addition to the standard capabilities that you would normally get from the open source Cacti package.

To ensure that all data are collected efficiently from the LSF environment, IBM provides specific data pollers (Platform data pollers) that work with IBM Platform RTM. You are free to use your own data pollers to work with RTM, but IBM Platform RTM does not provide support for those custom data pollers.

Chapter 2. Monitoring clusters with Platform RTM

You can use IBM Platform RTM to configure, monitor, and control LSF clusters.

Configure cluster management

Click the **Console** tab and use the left navigation pane to configure cluster management.

The left navigation pane has the following sections: Create (graphs), Management, License Management, Grid Management, Collection Methods, Templates, Presets, Import/Export, Syslog Settings, Configuration, and Utilities. Most are default Cacti utilities and features, and are documented in the Cacti documentation (<http://cacti.net/documentation.php>).

Perform the following actions to configure cluster management:

- Monitor your cluster by using Platform RTM.
Go to the **Management** section of the Console tab. For more information about management, see “Management section.”
- Add LSF clusters and complete certain database administration functions.
Go to the **Grid Management** section of the Console tab. For more information about grid management, see “Grid Management section” on page 6.
- Monitor license servers and pollers.
Go to the **License Management** section of the Console tab.
- Use Cacti utilities and features.
Go to the **Collection Methods**, **Templates**, or **Import/Export** sections of the Console tab to use Cacti utilities and features. For more information, see the Cacti documentation (<http://cacti.net/documentation.php>).

Management section

The **Management** section is in the Console tab.

Thresholds page

Go to the **Thresholds** page by clicking **Thresholds** under the Management section of the Console tab. Configured thresholds in your cluster will be displayed. A threshold triggers an alert if your clusters, hosts, queues, or jobs meet the conditions of the threshold.

- Name. The name of the cluster or host and the threshold. Click the name to change the threshold settings.
- Type. The type of threshold (for example, High/Low, Baseline, and Time Based)
- High. The high threshold boundary value. If the current value of the monitored data source item is greater than this value for a specified duration, the threshold triggers an alert.
- Low. The low threshold boundary value. If the current value of the monitored data source item is lower than this value for a specified duration, the threshold triggers an alert.
- Trigger. The amount of time that the data source item must be in breach of the threshold before the threshold triggers an alert.

- **Duration.** If the data source item is still in breach of the threshold, the duration is the amount of time from when the alert was first triggered.
- **Repeat.** The amount of time that the threshold waits before the alert repeats if the data source item is still in breach of the threshold.
- **Current.** The current value of the monitored data field.
- **Triggered.** Indicates whether this threshold triggered an alert
- **Enabled.** Indicates whether this threshold is active.
- **Ack.** Indicates whether the threshold alerts are acknowledged: "on" indicates that the threshold is acknowledged; "off" indicates that the threshold either is not acknowledged, or had its acknowledgement reset.

Threshold Item

Go to the **Threshold Item** page for a threshold by clicking the name of the threshold from the Thresholds page. Configure threshold settings and event triggering from the page.

Event triggering behavior is based on realert cycle settings. When the threshold first triggers an alert, the event trigger starts based on a high or low threshold breach. If the alert stays triggered, the event trigger is started again unless the realert cycle is set to **Never**. When the alert reverts to normal, the threshold triggers the norm threshold command or script.

You can configure the following items from here:

- **Template propagation enabled:** Enable the propagation of changes to the threshold template
- **Threshold name:** The name of the threshold as it is displayed in the Name column in the list of thresholds.

Note:

You can use placeholders to customize your threshold name. Placeholders for the threshold name are enclosed by pipe characters (|), for example, `|cluster_name|`.

- **Threshold enabled or disabled**
- **Weekend exemption:** Disable threshold alerts on weekends
- **Disable restoration email:** Disable threshold alerts when the threshold returns to normal
- **Reset acknowledgement:** Reset acknowledgements when the threshold returns to normal
- **High/low threshold values**
- **Threshold type:** High/low, baseline, or time based.
- **Event triggering (Shell command):** Specifies event trigger commands or shell scripts in the event of a breach.
 - **High Threshold Trigger Command/Script:** If the threshold is breached because the data source exceeds this value, the threshold triggers the specified command or shell script.
 - **Low Threshold Trigger Command/Script:** If the threshold is breached because the data source drops below this value, the threshold triggers the specified command or shell script.

- Norm Threshold Trigger Command/Script: If the threshold is breached, then returns to normal, the threshold triggers the specified command or shell script.
- Event triggering (Grid administrator host level triggers): Specifies host-level actions in the event of a breach.
 - Host Level Action (High Threshold): If the threshold is breached because the data source exceeds this value, the threshold triggers the specified action on the host.
 - Host Level Action (Low Threshold): If the threshold is breached because the data source drops below this value, the threshold triggers the specified action on the host.
- Email message body: Email alert message content. This control specifies the template that is used in alert email notifications for this threshold.

Note:

You can use placeholders to customize your alert emails and provide more information. Placeholders for the email message body are enclosed by angle brackets (<>), for example, <cluster_name>.

- Syslog settings
- Data type: Special formatting for the data
- Realert cycle: The amount of time the threshold repeats the alert, if it is still in breach.
- Notify accounts and extra alert emails: Email addresses to be notified when the threshold raises an alert

Placeholder tags

Placeholders are custom tags that represent real system values. You can insert placeholders in threshold names to show customized names that are based on your system and you can insert placeholders in alert email templates to present more information for administrators. Placeholders make it easy for them to follow up on the alert.

Tags for threshold names are enclosed by pipe characters (|), while tags for alert email templates are enclosed by angle brackets (<>). Not all placeholders are available for threshold names; some placeholders are only available for alert email templates. The following is a list of the placeholders available for your thresholds:

Table 1. Names, tags, and descriptions of available placeholders

Placeholder name	Tag for threshold name	Tag for alert email template	Description
Cluster ID	clusterid	<clusterid>	The ID of the cluster.
Cluster name	cluster_name	<cluster_name>	The name of the cluster.
Cluster LSF master	cluster_lsfmaster	<cluster_lsfmaster>	The name of the LSF master host for the cluster.
Cluster LSF version	cluster_version	<cluster_version>	The version of LSF running in the cluster.
Cluster LSF LIM port	cluster_limport	<cluster_limport>	The port number of LIM running in LSF on the master host.

Table 1. Names, tags, and descriptions of available placeholders (continued)

Placeholder name	Tag for threshold name	Tag for alert email template	Description
Custom data value	custom_custom_field_name	<custom_custom_field_name>	The custom data value from the data source that is linked in this alert. For example, custom_percent, custom_status.
Host name	host_hostname	<host_hostname>	The host name of the device that is linked in this alert.
Host description	host_description	<host_description>	The host description of the device that is linked in this alert.
Threshold description	Not available	<DESCRIPTION>	The threshold description.
Threshold host name	Not available	<HOSTNAME>	The host name of the threshold.
Threshold trigger time	Not available	<TIME>	The time in which the threshold triggered this alert.
Threshold graph URL	Not available	<URL>	The link to the URL of the threshold graph.
Threshold current value	Not available	<CURRENTVALUE>	The current value of the data field that is being monitored by the threshold, at the time of the alert email.
Threshold name	Not available	<NAME>	The name of the threshold.
Threshold data source name	Not available	<DSNAME>	The name of the data source that is being monitored by the threshold.
Threshold type	Not available	<THOLDTYPE>	The threshold type.
Threshold high value	Not available	<HI>	The high threshold boundary value.
Threshold low value	Not available	<LO>	The low threshold boundary value.
Threshold trigger	Not available	<TRIGGER>	The threshold trigger value.
Threshold graph ID	Not available	<GRAPHID>	The ID of the threshold graph.
Threshold duration	Not available	<DURATION>	The duration of the threshold.
Threshold details URL	Not available	<DETAILS_URL>	A URL to the threshold details page, which is a list of hosts that breached this threshold.
Threshold breached items	Not available	<BREACHED_ITEMS>	A list of items that breached this threshold, in an HTML table format.
Threshold graph	Not available	<GRAPH>	The threshold graph that is embedded into the email.
Threshold date	Not available	<DATE_RFC822>	The threshold date in RFC 822 format. For example, Thu, 01 Jan 2009 01:11:01 +0100

Grid Management section

The **Grid Management** section is in the Console tab.

Pollers page

Go to the **Pollers** page by clicking **Pollers** under the Grid Management section of the Console tab. Information about Platform RTM pollers is shown. These pollers collect information from the LSF cluster; Platform RTM uses this data to build various reports for Platform RTM users and administrators.

- **Add.** Add a poller. Click to open the **RTM Poller Edit [new]** page and specify the properties of the new poller.
- **Poller Name.** The defined name for the poller. Click a name to open the **RTM Poller Edit** page and edit poller properties (for example, the poller name, LSF version, bin directory location, poller location, and support information).
- **Poller ID.** The ID assigned to the poller.
- **Status.** The status of the poller. Statuses can be Up, Down, or N/A.
- **LSF Version.** The LSF version that is running on the associated cluster.
- **Physical Location.** The physical directory location of the local Platform RTM poller (for example, `/opt/rtm/lsfversion/bin`). If the directory is found and verified, the message [OK: DIR FOUND] is displayed below this field.
- **Support Information.** Enter a text string description for the location of the data poller (for example, "Data Center").
- **Choose an action.** Select one or more checkboxes to manage the pollers (for example, **Delete**), then select an action and click **go**.

Clusters page

Go to the **Cluster** page by clicking **Clusters** under the Grid Management section of the Console tab. Information about LSF clusters (including configured time-out thresholds, and job efficiency information) is shown and the pollers that collect data from them.

For more information about configuring a cluster with Platform RTM, see "Add or edit LSF clusters for Platform RTM to monitor" on page 17.

- **Add.** Add a cluster for Platform RTM to monitor. Click to open the **Cluster Edit [new]** page and specify the properties of the cluster.
- **Cluster Name.** The defined name for the cluster. Click a name to open the **Cluster Edit** page and edit cluster properties, defaults, and various collection settings.
- **Cluster ID.** The ID assigned to the cluster.
- **Poller Name.** The name of the poller that is associated with this cluster.
- **Collect Status.** The current data collection status for this cluster. Status can be Disabled, Up, Jobs Down, Down, Diminished, Admin Down, and Maintenance.
 - **Maintenance.** Indicates that RTM is conducting Database maintenance.
 - **Admin Down.** This status is displayed if you cleared the **Should Daemons be Enabled** option in the General Settings tab for the Poller.
 - **Up.** Indicates that there are queues, hosts, and loads that are collected within 3*non-job-interval and job information that is collected within 3*job-interval. Options are set in the **Max Allowed Runtime** for Queue/Host/Load Collection Settings and Job Collection Settings for the Poller.
 - **Down.** Indicates that there are no queues, hosts, and loads that are collected within 3*non-job-interval and job information collected within 3* job-interval.
 - **Jobs Down.** Indicates that only the job information was not collected within 3*job-interval.

- Diminished. Indicates that: One or several of queues, jobs, hosts, and loads information are collected but not all of them are collected within the related intervals. Jobs collection is related to job-interval while queues/hosts/loads information are related to non-job-interval.

This status can also indicate that the collection of jobs, queues, hosts, and loads information never started.

If you see No data shown above each status indicates no data.

- Efic Status. An indicator of job efficiency within this cluster, based on configured thresholds. Status can be OK, Recovering, Warn, Alarm, and N/A. Thresholds are set from **Console > Grid Settings > Status/Events**.
- Efic Percent. An indicator of the average efficiency of running jobs within the cluster, reported as a percentage. The minimum runtime setting can be set from **Console > Grid Settings > Status/Events**.
- Total Hosts. The total number of hosts in this cluster.
- Total CPUs. The total number of processors in this cluster.
- Total Clients. The total number of clients in this cluster.
- Collect Freq. The configured data collection frequency.
- Collect Timeout. The configured data collection timeout.
- Job Minor Freq. The configured job minor frequency.
- Job Major Freq. The configured job major frequency.
- Job Timeout. The configured job timeout.
- LIM Timeout. The configured lim timeout.
- Choose an action. Select one or more check boxes to manage the clusters (for example, **Enable** or **Disable**), then select an action and click **go**.

Utilities page

Go to the **Utilities** page by clicking **Grid Utilities** under the Utilities section of the Console tab. Information about Platform RTM utilities as related to database administration (such as data backup, purging, and record removal), along with status information about cluster pollers is shown.

- View Grid Process Status. Click to open the **Grid Process Status** page and show status information that is associated with cluster polling processes (for example, statistics for cluster poller, runtime, database maintenance, license collection).
- Force Cacti Backup. Click to back up key Cacti and Platform RTM database tables. For more information about backup and restore, see “Maintaining Database” on page 70.
-
- Manage Grid Hosts. Click to open the **Manage Hosts** page and selectively remove client records from the host database.
- Backup Files. Click a file name to download the backup file. The section is not displayed until Platform RTM creates at least one backup file.

Configure cluster interaction

Click the **Console** tab and refer to the Console tab on the left side to configure cluster interaction.

The Console tab has 11 sections: Create (graphs), Management, License Management, Grid Management, Collection Methods, Templates, Presets,

Import/Export, Syslog Settings, Configuration, and Utilities. Most are default Cacti utilities and features, and are documented in the Cacti documentation (<http://cacti.net/documentation.php>).

To configure cluster interaction with Platform RTM do these steps:

- Configure the LSF clusters to interact with Platform RTM.
Click **Grid Settings** under the **Configuration** section of the Console tab. For more information about the subtabs, see “Grid Settings page.”
- Configure Platform RTM to interact with LSF clusters.
Click **Settings** under the **Configuration** section of the Console tab. For more information about the subtabs, see “Cacti Settings page” on page 12.

Grid Settings page

Go to the **Grid Settings** page by clicking **Grid Settings** under the Configuration section of the Console tab.

There is a subtab for each category of grid settings that you can change.

- **General:** Configure the default user settings for your cluster.
Among the fields that are displayed, take note of the following specific fields:
 - Strip Domains
Makes the name output shorter for hosts that belong to the specified domain names. You can conserve display space if you are displaying hosts in common domains within your organization.
 - Summary Host name Substitute
Conserve display space for common host name substrings.
 - Minimum User Screen Refresh Interval
Restricts the minimum refresh interval that your users can set to reduce load on the system. Set a higher refresh interval for larger clusters to reduce system load.
 - Maximum Job Zoom Time Range
Restricts the maximum job time in which your users can zoom in after which they cannot zoom in anymore. Restrict this setting to reduce load because the job zoom function is system-intensive. Set a smaller window for larger clusters to reduce system load.
 - User Group Filter Operation
Specifies how your cluster handles user group filtering. There are two ways to filter jobs by user groups. Setting the user group filter operation to **User Group Membership** displays all the jobs of the user that belong to the group. The **Job Specification** option displays only the jobs that are submitted with **bsub -G** option.

To set the method that matches your configuration:

1. Click the **Console** tab.
2. Under the Configuration section of the Console tab, click **Grid Settings**.
3. Click General tab.
4. Select the User Group Filter Operation (**User Group Membership** or **Job Specification**) from the list.
5. Click **Save**.

To see the display, go to the **Grid** tab then click **Groups** under User/Group Info section of the Grid tab. Click the Group Name link to view the list of jobs for the group based on the specified setting.

- Maximum Export Rows
Restricts the maximum number of rows that your users can export to improve system performance. Increasing the maximum number of rows that are exported reduces system performance during export operations.
- Cluster CPU Factor Leveling

Important:

Do not enable this setting unless you understand how to apply CPU factoring to hosts in your cluster. If you want to learn more about CPU factors, then refer to the *Administer LSF* documentation.

- Multiple Column Sort
Sets multiple column sorting. The default number of columns is 1 and the maximum number of sorting columns is 4. This global setting can be configured by selecting the number of columns on the **Multiple Column Sort** list.

Note:

- Upon initial access to a page, a preselected column for sorting is displayed. The display varies depending on the page you are viewing.
To add a sort column, click the column name. The first column that you selected is automatically assigned with the first order sorting, and so on.
When you reach the maximum number of sorting columns and you continue to add a next column, the last order sorting is replaced with the new column.
- The column sort order is based on the sequence of columns selected. A number beside the column name represents the sorting priority. The arrow represents the direction of sort (starts with 1).
- You can adjust the sorting direction by clicking the column name to change the direction of sorting. The priority order of the sorting columns does not change.
- To delete a sorting column, click the x beside the column name. After a sorting column is deleted, the rest of the columns automatically shift upward.
- License and Grid Plug-in pages now support multiple column sort capability.
- All Platform RTM pages support the multiple column feature except for pages under the Console tab: Grid Management and Collection Methods list

Important:

- Multiple Column Sorting is controlled by realm permissions and is enabled on a per user basis in user management. This setting is disabled by default.
To enable realm permission, go to **Console** tab then click **User Management** under Utilities section of the Console tab. Select the **User** then click **Realm Permissions** tab. Click the **RTM -> Multiple Column Sort** check box.
- Enable Ajax filter. Enable this feature to improve your search criteria. Enter at least three '%' characters in the Grid tab to get a list of all hosts, host group, users, or user group.
- **Poller:** Configure poller defaults for data collection, interval settings, and thresholds.

- **Maint:** Configure system maintenance settings. You can keep more data for smaller clusters, because there are fewer records for these clusters.

Among the fields that the page displays, take note of the following specific fields:

- Detail Job Data Retention Period

Records job details such as total process count, threads, memory, PID, PGID, and swap, and these details are used to generate Job Graphs.

The size of each job record depends on job volume and your cluster settings. The system can hold a maximum of 10 million records. Use this upper limit along with the approximate number of jobs per week in your cluster to determine the ideal retention period.

- Summary Job Data Retention Period

Individual job records are kept for this period after the job ended. The size of each job record depends on job volume and your cluster settings.

- Daily Statistics Retention Period

Records of daily summary statistics are kept for this period after the job ended. As these records are added every day, you can keep records for a longer time, depending on the job volume. Smaller clusters with less than one million jobs per year can have a retention period as high as three years.

- Backup Cacti Database

This field enables a disaster recovery backup to restore your Cacti and RTM configuration. Some job data is lost during the database restoration, though you can use other utilities to restore all the job data.

Note:

Database backup files are disk-intensive for larger clusters.

- **Archiving:** Configure database archiving settings.

Data archiving allows you to send job data to a non-local MySQL database to be used by other tools such as Cognos, Crystal Reports, and Tableau. Currently, only MySQL is supported as an external database. However, if you want to add another database, then contact Support.

- **Paths:** Configures cluster directories and file paths. If the directories and file paths are found and verified, the message [OK: FILE FOUND] displays below the corresponding fields.

- **Thold:** Configures cluster alert and threshold settings, including alerts and thresholds to identify when resources are idle/closed, low, busy, or starved.

You can also set how the RES-Down Status is treated when you search for jobs By Host to By Group. On the Down Service Status section, select Out of Service or In Service. The specified option becomes available in the Status field on the Batch Host Filters when you click **By Host** or **By Host Group** under the Job Info section of the Grid tab. If you choose In Service, jobs with RES-Down status are not shown on the list. To see jobs with RES-Down status, choose Out of Service.

- **Aggregation:** Configures default behavior for project information aggregation, host information aggregation, and memory tracking.

Among the fields that the page displays, take note of the following specific fields:

- Wall clock Calculation Method

Set this field for chargeback calculations, depending on whether you charge for suspend time.

- Should Project Names be Aggregated?

Enable this field to collect job data based on project names.

- Project Aggregation Method

Project aggregation is used for project names that contain hierarchical metadata to assist with tracking.

- Should License Project Job Performance be Tracked?

Enable this field to track running jobs that are associated with a License Project in **Grid** -> **By License Project**. By default, it is disabled and license project information is not shown.

- Aggregation Level

When you build a job group hierarchy, only a number of fields are significant for analysis. RTM aggregates job groups by only the specified level, if you set up a positive integer value.

- **Status/Events:** Configure default behavior, thresholds, and visual cues for job flapping, cluster and job efficiency, PID levels, and job dependencies. If you set the color to **None**, the legend items does not appear, and colorizing of the GUI will not happen on the **Job Details** page.
- **Idle Jobs:** Configure detection and email notification settings for idle jobs. You can configure the following global settings: Filter Name, Email Subject, Email Message, and Legend Background Color. Specific Idle job settings are set under the **Cluster Management** page in the **Idle Job**stab.
- **Memory Exceptions:** Configures detection and email notification settings for memory usage violations.
- **Runlimit Exceptions:** Configures detection and email notifications for job runtime limit exceptions.

Cacti Settings page

Go to the **Cacti Settings** page by clicking **Settings** under the Configuration section of the Console tab. There is a subtab for each category of Cacti settings that you can change.

If you need more details about any of the Cacti features, refer to the Cacti documentation (<http://cacti.net/documentation.php>).

- **General:** Configure the default Cacti settings.
- **Paths:** Configure Cacti directories and file paths. If the directories and file paths are found and verified, the message [OK: FILE FOUND] displays below the corresponding fields.
- **Poller:** Configure default poller settings.
- **Graph Export:** Configure graph export settings.
- **Visual:** Configure Cacti display settings.
- **Authentication:** Configure Cacti authentication settings.
- **Boost:** Configure Boost server settings.
- **Mail/DNS:** Configure email and DNS settings for the Cacti server.

Among the fields that the page displays, take note of the following specific field:

- **Sendmail Path.** Specify the location of the **sendmail** binary file on the RTM host, if sendmail is selected as the mail service. If the file is found and verified, the message [OK: FILE FOUND] displays below this field.
- **Misc:** Configure third-party plug-ins. You can also view Cacti Log.
- **Alerting/Thold:** Configure alert and cluster threshold settings.
- **Syslog:** Configure Syslog event settings.

- **RTM Plugins:** These plug-ins belongs to RTM. Configure Platform RTM plug-in settings.

Among the fields that the page displays, take note of the following specific fields:

- DNS suffix for the RTM server. This setting is found in the **ssh** or **telnet** configuration of the RTM host and is only required if your web browser cannot resolve host names.
- Add All Features Graph. If this setting is enabled, RTM produces all graphs. Since there are many graphs, this process takes significant amount of time and disk space. It reruns every 24 hours to add new graphs if new features are available.

When disabled (cleared), only the FLEXlm License - Feature Use and FLEXlm License - License Server Usage graphs are produced.

- Exactly Filter License Checkout. If this setting is enabled, RTM filters license checkout only by host, user, vendor, and feature name when forwarded to License Checkout. Enable this setting for larger systems where there are several FLEXlm License Services are being monitored. This increases performance. Make sure that the DNS information in FLEXlm matches the LSF hostname information.

Control an LSF cluster

Platform RTM can control LSF clusters, hosts, queues, and jobs, if you enable RTM to control the LSF clusters. RTM controls the LSF clusters by starting LSF commands in the LSF master host.

- Cluster-level components

You can control the following cluster-level components:

- **mbatchd** (start, restart, or shut down)
- **LIM** (start, restart, or shut down)
- **RES** (start, restart, or shut down)

You can also run the following LSF commands:

- **badmin reconfig** (dynamically reconfigures LSF)
- **lsadmin reconfig** (restarts LIM on all hosts in the cluster)

- Hosts

You can open or close hosts in LSF clusters.

- Queues

You can do the following actions to control queues in an LSF cluster:

- open queues
- close queues
- activate queues
- deactivate queues
- switch all jobs from one queue to another

- Jobs

You can run the following LSF commands to control jobs in an LSF cluster:

- **btop** (moves a pending job relative to the first job in the queue)
- **bbot** (moves a pending job relative to the last job in the queue)
- **bswitch** (switches unfinished jobs from one queue to another)
- **bresume** (resumes suspended jobs)
- **brun** (forces jobs to run immediately)

- **bstop** (suspends unfinished jobs)
- **bkill** (sends signals to kill unfinished jobs)
- Force kill (forces a job kill by using the **bkill** **-r** command)
- Signal kill (sends a specific signal to kill a job by using the **bkill** **-s** command)

Run grid control commands on an LSF cluster

Run LSF commands on an LSF cluster through the Platform RTM Console.

Before you begin

You must ensure that you enabled grid control on the LSF cluster.

Procedure

1. From the Platform RTM console, click the **Console** tab.
2. Under the Grid Management section of the Console menu bar, click **Clusters**.
3. Click the Cluster Name link of the cluster you want to configure.
4. Configure the cluster settings. If the grid control requires more information, specify the details in the displayed fields.
5. Click **Save**.

Results

To see the status of a cluster, refer to “Dashboards section” on page 21.

Enforce cluster action control comments

Platform RTM can enforce mandatory cluster action control comments for the following:

- By Host
- By Queue
- Cluster

You must set the **Mandatory Cluster Action Control Comments** option to allow your Grid Administrator to enter comments before a control action (such as startup/restart/shutdown, open/close) for host, queue, or cluster can be submitted.

Set mandatory cluster action control comments

Procedure

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings**.
3. Click **General** tab.
4. Under Cluster Control Settings, check the **Mandatory Cluster Action Control Comments** check box.
5. Click **Save**.

Comment on cluster action control by host, queue, or cluster

Procedure

1. Click the **Grid** tab.

2. Under the Job Info section of the Grid menu bar, click **By Host**, **By Queue**, or **By Cluster**.
3. Check one or more boxes corresponding to the Host Names you want to act on.
4. Choose an action from the list.
5. Click **Go**.
A confirmation dialog box is displayed.
6. Enter your comments in the text box.
If the mandatory cluster action comment was set, you must enter your comments to proceed with the submission.
7. Click **Yes**.

Configure idle job detection for cluster

About this task

Idle jobs are configured per cluster on the Cluster Edit page and they are not configured globally. Specific Queues might be excluded in idle job calculation, such as jobs submitted to an interactive Queue.

Note:

After a fresh installation of RTM 9.1, any previous settings for idle jobs are no longer in effect. You must configure them again.

Procedure

1. Click the **Console** tab.
2. Under the Grid Management section of the Console menu bar, click **Clusters**.
3. Click the Cluster Name of the cluster for which you want to set the idle job detection.
4. Check the **Enable Idle Job Detection** box if you want to search for idle jobs in the cluster.
5. You can choose to set the values for the following fields:
 - **Email Notification Type**
 - **Minimum Runtime**
 - **Floating Window**
 - **CPU Time Threshold**
 - **Include Job Types**
 - **Job Commands**
 - **Idle Jobs Exclude Queues**
6. Click **Save**.

Configure idle job detection (Grid settings)

Procedure

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings** and go to **Idle jobs**.
3. Set the values for the Idle Job Detection fields:
 - **Filter Name**
Provide the filter name to show the legend and Job exception.

- **Email Subject**
Define the email subject for the idle job email message by using these replacement tags: <JOBID>, <INDEXID>, <SUBMITTIME>, <CLUSTERNAME>, <USER>, and <CPUSECS>
- **Email Message**
Define the email content by using these replacement tags: <JOBID>, <SUBMITTIME>, <CLUSTERNAME>, <USER>, and <CPUSECS>
- **Legend Background Color**
Define the legend. This feature is disabled if Color is set to None.

4. Click **Save**.

Configure memory violation detection

About this task

You can configure the memory settings for jobs to help you monitor the memory usage and flexibly control jobs.

Procedure

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings** and go to **Memory Exceptions**.
3. Set the values for the Memory RUSAGE Violations fields:
 - **Enable Memory RUSAGE Job Detection**
Search for memory violation jobs.
 - **Email Summary Reports**
Account to whom the email summary report notification is sent.
 - **Email Schedule**
Notification frequency for receiving email. Emails are sent weekly on Sundays during database maintenance.
 - **Memory Overusage Filter Name**
Filter name to be displayed in the job details legend for memory overuse.
 - **Memory Overusage Allocation**
Memory usage percentage above the RUSAGE level that is acceptable before the job is flagged.
 - **Memory Overusage Background Color**
Color of the legend. This feature is disabled if Color is set to None.
 - **Memory Underusage Filter Name**
Filter name to be displayed in the legend and in the Job exception filter display.
 - **Memory Underusage Allocation**
Memory usage percentage below the RUsage level that is acceptable before the job is flagged.
 - **Memory Underusage Background Color**
Color of the legend and row display. This feature is disabled if the color is set to none.
 - **Minimum Run Window**
Minimum run time.

- **Minimum Memory Limit**

Minimum memory limit for your jobs.

- **Email Subject**

Defined with these replacement tags: <OVERFILTER>/<UNDERFILTER> and <CLUSTERNAME>

- **Email Message**

Defined with these replacement tags: <CLUSTERNAME>, <OVERFILTER>, <OVERSHOOT>, <UNDERFILTER>, <UNDERSHOOT>, and <REPORTTABLE>

4. Click **Save**.

Add or edit LSF clusters for Platform RTM to monitor

The following is a list of actions that you must take, in chronological order, to add or edit LSF clusters for Platform RTM to monitor:

1. “Add or edit clusters”
2. “Add the RTM host to the LSF cluster as an LSF client” on page 19
3. “Add all hosts in the LSF cluster to RTM as devices” on page 20

Add or edit clusters

About this task

Add or edit any LSF clusters that you want Platform RTM to monitor by either of the following methods:

Procedure

- “Add or edit clusters by using the Platform RTM Console”
You can use the Platform RTM Console to add or edit an LSF cluster for monitoring.
- “Add clusters to RTM by using a script” on page 18
You can use a script to add an LSF cluster for monitoring by Platform RTM.

Add or edit clusters by using the Platform RTM Console

About this task

Use the Platform RTM Console to add an LSF cluster for monitoring by Platform RTM.

Procedure

1. Click the **Console** tab.
2. Under the Grid Management section of the Console menu bar, click **Clusters**.
3. Choose to add or edit a cluster.
 - To add an LSF cluster, click **Add**.
 - To edit an existing LSF cluster, click the name of the cluster that you want to edit.
4. Specify (or update) the required fields that describe your LSF cluster.
At a minimum, you must specify the following fields to add the LSF cluster:
Cluster Name, **LSF Master LIM Hostname**, **LSF Master LIM Port**, **Grid Poller**, and **Primary LSF Administrator Username**.

If you want to monitor multiple clusters and if the Platform RTM host is an LSF server, make sure the clusters use different LIM Ports. However, it is recommended to install Platform RTM on the LSF client host.

For the **Grid Poller** field, select the appropriate poller for your version of the LSF cluster.

In this release, you can now select Poller for LSF 8 from the list of **Grid Poller**.

5. Click **Create** (or **Save**) to save the settings for your LSF cluster.

What to do next

- If you edited an LSF cluster that was already in the Platform RTM Console, you do not have to do anything else.
- If you added an LSF cluster to the Platform RTM Console, you must add the Platform RTM host to the LSF cluster as described in “Add the RTM host to the LSF cluster as an LSF client” on page 19.

Add clusters to RTM by using a script

About this task

Use the `grid_add_cluster.php` script to add an LSF cluster to Platform RTM.

Procedure

1. From the command line, go to the `plugins/grid` subdirectory of the Cacti installation directory.
2. Use **php** to run the `grid_add_cluster.php` script.

```
php -q grid_add_cluster.php --type=0 --pollerid=lsf_type  
--cluster_name=cluster_name_text --cluster_env=lsf_envdir_path
```

where

- *lsf_type* is an integer that represents the version of LSF running in the cluster:

- 1: LSF 6.1
- 2: LSF 6.2
- 3: LSF 7.0.1
- 4: LSF 7.0.2
- 5: LSF 7.0.3
- 6: LSF 7.0.4
- 7: LSF 7.0.5
- 8: LSF 7.0.6
- 9: LSF 8.0

Note:

The Poller ID does not necessarily follow an order. You can add your own pollers in a different order. It is best to use the Platform RTM Console to identify the Poller ID you want to use.

Poller IDs are maintained during upgrade and only new pollers are inserted into the database.

- Go to the **Console** tab then click **Pollers** under Grid Management section. The RTM Pollers page opens displaying the Poller Names and Poller ID. Use the Poller ID value in the list for the `grid_add_clusters.php`.
- *cluster_name_text* is the name of the cluster

- `lsf_envdir_path` is the path to the `lsf.conf` file for your LSF cluster.
- For example, to add an LSF 7.0.6 cluster named `maincluster` with `lsf.conf` in `/share/lsf/conf`:
- ```
php -q grid_add_cluster.php --type=0 --pollerid=8
--cluster_name=maincluster --cluster_env=/share/lsf/conf
```

## What to do next

If you add clusters to RTM by using a script, you must:

1. Add the Platform RTM host to the LSF cluster as described in “Add the RTM host to the LSF cluster as an LSF client.”
2. Verify that the new cluster is added to RTM from the RTM Console by clicking **Clusters** in the Console menu bar and checking that the new cluster is up.

## Add the RTM host to the LSF cluster as an LSF client

### About this task

For any LSF cluster that RTM monitors, you must add the RTM host to the cluster as an LSF client to give RTM access to LSF cluster data.

### Procedure

1. Log in to the LSF master host.
2. If the LSF master host cannot resolve the Platform RTM host name to an IP address, edit the `/etc/hosts` file and add the IP address and host name of your RTM host.

#### Note:

If you can successfully ping the Platform RTM host name from the LSF master host, you can skip this step.

3. Edit the `lsf.cluster.cluster_name` file and add the RTM host to the Host section.
4. Reconfigure LIM and restart `mbatchd` to apply your changes to the cluster.
 

```
lsadmin reconfig
badmin mbdrestart
```
5. Test that you added the RTM host successfully to the LSF cluster.
  - a. Log in to the RTM host.
  - b. From the RTM host, use **telnet** to log in to the LSF LIM port of your RTM host.

The default LIM port is 6879 for LSF 6.2 clusters and 7869 for LSF 7.x clusters.

For example, for LSF 7.x clusters,

```
telnet<LSF Master IP>7869
```

If you connect to the IP address of the LSF master host, you added the RTM host successfully.

## Results

After you add the Platform RTM host to the LSF cluster, Platform RTM can now monitor the LSF cluster.

## What to do next

If you want Platform RTM to be able to monitor all individual hosts in the LSF cluster, add all hosts in the LSF cluster to Platform RTM as described in “Add all hosts in the LSF cluster to RTM as devices.”

## Add all hosts in the LSF cluster to RTM as devices

### About this task

Optional. Use the `grid_add_cluster.php` script to add all hosts in the cluster list to the device list in RTM. This addition enables Platform RTM to monitor the hosts in the cluster.

### Procedure

1. In the command line, go to the `plugins/grid` subdirectory of the cacti installation directory.  
For example, `cd /opt/cacti/plugins/grid`
2. Use **php** to run the `grid_add_cluster.php` script to add all hosts in the cluster list to RTM as devices and create related graphs by using the Grid Host template.  

```
php -q grid_add_cluster.php --type=1 --clusterid=all --template=14
```

### Results

You can view these hosts in the device list by clicking **Devices** in the Management section of the Console menu bar.

---

## View LSF cluster and job information

Click the **Grid** tab and refer to the Grid menu bar on the left side to view LSF cluster and job information.

For more information about configuring the interaction between Platform RTM and LSF, see “Configure cluster interaction” on page 8.

The following is a list of actions that you can take to view LSF cluster and job information.

- Review a summary of cluster and host health.  
Go to the **Dashboards** section of the Grid menu bar. For more information about the pages in the section, see “Dashboards section” on page 21.
- Control the cluster.  
Go to the **Management** section of the Grid menu bar.  
The pages can be used to control your cluster, and are only available to users with the **Cluster Control Management** realm permission (Click **User Management** under the Utilities section of the Console menu bar).
- Review information about LSF jobs.  
Go to the **Job Info** section of the Grid menu bar.  
The pages provide information about LSF jobs, either at the level of host, host group, queue, job array, application, or group. You can also view detailed information about specific jobs and pending reason history of jobs. For more information about the pages in the section, see “Job Info section” on page 23.
- Review information about LSF users and user groups.



Go to the **User/Group Info** section of the Grid menu bar. For more information about the pages in the section, see “User/Group Info section” on page 33.

- Review information about the load on hosts and host groups.

Go to the **Load Info** section of the Grid menu bar. For more information about the pages in the section, see “Load Info section” on page 34.

- Review information about LSF hosts and host groups.

Go to the **Host Info** section of the Grid menu bar. For more information about the pages in the section, see “Host Info section” on page 35.

- Review daily statistics, FLEXlm license usage, and job parameters.

Go to the **Reports** section of the Grid menu bar. You can also filter statistics and batch system parameters for specific information from here. For more information about the pages in the section, see “Reports section” on page 37.

- Configure the **Grid** tab settings.

Click the **Grid** tab, then the **Settings** tab on the right side of the tabbed interface to configure items in the **Grid** tab. For more information, see “Settings tab (from the Grid page)” on page 38.

## Dashboards section

The Dashboard section is in the Grid menu bar.

Together, the RTM dashboards display useful information about the status of your LSF clusters. By changing the icon color, RTM can also alert operators when a host becomes unavailable for some reason.

In its current form, you can view the status of each of your clusters, the status of feature licenses, and a pictorial representation of the hosts on those clusters. If you choose to filter the display, the display is changed to reflect the current filtering.

### Reasons Dashboard

The reasons dashboard shows pending reason statistics for the cluster grouped by project or queue.

### Statistical Dashboard

The statistical dashboard presents non-time series cluster and host statistics by using graphically rich charts.

### Cluster Dashboard

The cluster dashboard shows the following information:

- Cluster Name: The LSF cluster name.
- Cluster Status: The status of the cluster.
- Master Status: The status of the master host in the cluster.
- PAU: The type of the host currently controlling the cluster. Valid values are as follows:
  - P: Primary master host
  - A: Failover host
  - U: Unknown host type
- Collect Status: The data collection status for the cluster.
- CPU %: The cluster’s overall CPU utilization rate, as a percentage.

- Slot %: The entire cluster's slot utilization, as a percentage.
- Efic %: The entire cluster's CPU efficiency for running jobs. Efficiency is calculated with this formula:  $cpu\_time / (run\_time \times \#\_of\_cpus)$ .
- Total CPUs: The total number of CPUs in the cluster.
- Host Slots: The total number of slots available to run jobs in the cluster.
- Pend Jobs: The total number of pending jobs in the cluster.
- Run Jobs: The total number of running jobs in the cluster.
- Susp Jobs: The total number of suspended jobs in the cluster (including system suspended and user suspended jobs).
- Hourly Started: The total number of jobs that are started during the last hour.
- Hourly Done: The total number of jobs that are completed during the last hour.
- Hourly Exit: The total number of jobs that are cancelled during the last hour (unsuccessful completion).

## Cluster Details page

Select the **Details** box in the Cluster dashboard to see a summary of status details of your clusters.

Status information that is represented by icons for LIM status and batch status is shown. Table 2 shows the icons that are used to display LIM status and their meanings.

Table 2. LIM status icons





| Icon                                                                                | Status                                                                                                                                            |
|-------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <ul style="list-style-type: none"> <li>• ok</li> </ul>                                                                                            |
|  | <ul style="list-style-type: none"> <li>• locked</li> <li>• locked-user</li> <li>• locked-window</li> <li>• locked-master</li> </ul>               |
|  | <ul style="list-style-type: none"> <li>• busy</li> </ul>                                                                                          |
|  | <ul style="list-style-type: none"> <li>• unavail</li> <li>• unlicensed</li> <li>• <b>sbatchd</b> is down</li> <li>• <b>RES</b> is down</li> </ul> |

Table 3 shows the icons that are used to display batch status and their meanings.

Table 3. Batch status icons




| Icon                                                                                | Status                                                                                        |
|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
|  | <ul style="list-style-type: none"> <li>• ok</li> <li>• Any closed except for admin</li> </ul> |
|  | <ul style="list-style-type: none"> <li>• closed-admin</li> </ul>                              |

Table 3. Batch status icons (continued)

| Icon                                                                              | Status                                                                                               |
|-----------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
|  | <ul style="list-style-type: none"> <li>• unavail</li> <li>• unlicensed</li> <li>• unreach</li> </ul> |

For a description of the RTM Status icons, open the Host Dashboard (click **Host** in the Dashboards section of the Grid menu bar) and view the Host Status Legend.

## License Dashboard

The license link on the dashboard takes you to the license plug-in that displays a list of your configured license servers.

## Host Dashboard

If you roll your mouse over a host, summary information displays about that host. For example, you can view load averages, numbers of job slots and current slot utilization, administrative notes and status. If you click a host icon, you are directed to the “RUNNING” jobs for that host (on the Job Info > Details page). Color-coding for the host icons is described under the Host Status Legend section.

The host icons can be displayed as either small or large. Click the **Settings** tab and modify the settings that are found under the **Visual** subtab to control this behavior.

The Host dashboard can notify you if there are any events triggered on the hosts. You can enable audible and/or visual notifications from the **Settings>General** tab. When you point to a host icon, you can see the **Alarm** tab, which lists any triggered alarms or grid alerts for that host. The Host dashboard shows alerts for Grid Alarm, Syslog, Threshold, and batchload. If required, you can stop, resume, or edit the selected alert.

## Alerts Dashboard

The alerts dashboard shows grid alerts. If required, you can stop, resume and edit the selected alert.

## Job Info section

The Job Info section is in the Grid menu bar.

### Host Job Statistics Viewer

Go to the **Host Job Statistics Viewer** by clicking **By Host** under the Job Info section of the Grid menu bar. Information about hosts in a cluster is shown.

- Host Name. The name of the host. Click a host name to show running jobs for this host (on the **Job Info > Details** page).
- Cluster. The LSF cluster to which this host belongs.
- Type. The type of host, as defined in the LSF configuration.
- Model. The model of the host, as defined in the LSF configuration.
- Load/Batch. The current Load and Batch status of the host.

If you choose Out of Service option on the Status field, jobs with RES-Down status are seen on this column. If you choose In Service, jobs with RES-Down status are not displayed here. You must set how the Down Service Satus is treated (through **Console > THold > Grid Settings** tab under Configuration list) for the corresponding option to be available in the Status filter.

- CPU Fact. The processor factor of the host, as defined in the LSF configuration.
- CPU Pct. The current processor utilization on the host.
- RunQ 1m. The exponentially averaged effective processor run queue length for this host over the last minute.
- Mem Usage. The percentage of memory usage of all jobs that are running on this host as a percentage of total memory.
- Page Usage. The percentage of page usage of all jobs that are running on this host as a percentage of total page size.
- Page Rate. The memory page scan rate of the host.
- Max Slots. The maximum number of job slots that can be allocated to this host.
- Num Slots. The number of jobs slots that are used by jobs that are dispatched to this host.
- Run Slots. The number of job slots that are used by jobs that are running on this host.
- SSUSP Slots. The number of job slots that are used by system-suspended jobs on the host.
- USUSP Slots. The number of job slots that are used by user-suspended jobs on the host.
- Reserve Slots. The number of jobs slots that are used by pending jobs that have job slots that are reserved within the host.

If graphs are created for this host, a graph icon is displayed to the left of the host name. Click the icon to view graphs for the host.

## Host Group Job Statistics

Go to the **Host Group Job Statistics Viewer** by clicking **By Host Group** under the Job Info section of the Grid menu bar. In many respects, the information that is shown is similar to that obtained through the LSF **bhosts** command with condensed host groups.

The Status filter is populated with all unique Load and Batch statuses that are currently experienced by hosts in any cluster.

Job information by LSF host group is shown:

- Host Group. The name of the LSF host group. Click a host group name to show running jobs for this group (on the **Job Info > Details** page).
- Cluster. The LSF cluster to which this host group belongs.
- Load/Batch. The current Load and Batch status for the host group. If no Status filter is set, this field shows N/A. Otherwise, it shows the current value that is selected for the Status filter.

If you choose Out of Service option on the Status field, jobs with RES-Down status are displayed. If you choose In Service, jobs with RES-Down status are not displayed. You must set how the Down Service Satus is treated (through **Console > THold > Grid Settings** tab under Configuration list) for the corresponding option to be available in the Status filter.

- Total Hosts. The total number of hosts in this host group.

- Avg CPU %. The average processor utilization for hosts in this host group.
- Avg r1m. The average exponentially averaged effective processor run queue length for this host group over the last minute.
- Avg Effic. The average efficiency of the host group.
- Total CPU. The overall processor utilization rate of the host group.
- Max Memory. The maximum memory that is used by the host group.
- Max Swap. The maximum swap usage of the host group.
- Max Slots. The maximum number of job slots available for this host group.
- Num Slots. The number of jobs slots that are used by jobs that are dispatched to this host group.
- Run Slots. The number of job slots that are used by jobs that are running on this host group.
- SSUSP Slots. The number of job slots that are used by system suspended jobs on the host group.
- USUSP Slots. The number of job slots that are used by user suspended jobs on the host group.
- Reserve Slots. The number of jobs slots that are used by pending jobs that have job slots that are reserved within the host group.

You can view Job Statistics graphs by clicking the corresponding **View Host Group Graphs** action icon of the Host Group. These graphs present images and statistics to help you control jobs flexibly.

- GRID - Host Group - CPU Utilization  
Presents statistics on processor utilization when the host is the server and its status is not in Unavail, Unreach, Unlicensed, or Closed-Admin.
- GRID - Host Group - Available Memory  
Presents a general image of available memory by host group level for more flexibility in controlling jobs.
- GRID - Host Group - Host Details  
Presents a general image of memory usage by host group level.
- GRID - Host Group - Memory Stats  
Presents statistics on average or maximum memory when job status is RUNNING, USUSP, and SSUSP.

## Project Viewer

Go to the **Project Viewer** by clicking **By Project** under the Job Info section of the Grid menu bar. Resources are shown in a cluster by project.

The information that is shown on the page is as follows:

- Project Name. The name of the project. Click a project name to show running jobs for this project (on the **Job Info > Details** page).
- Cluster Name. The LSF cluster to which this project belongs.
- Total Slots. The total number of job slots that are used for this project.
- Pending Slots. The number of job slots that are used by pending jobs for this project.
- Running Slots. The number of job slots that are used by running jobs for this project.
- Avg Effic. The average efficiency of this project.

- Max Mem. The maximum memory that is used by this project.
- Avg Mem. The average amount of memory that is used by this project.
- Max Swap. The maximum swap space that is used by this project.
- Avg Swap. The average swap space that is used by this project.
- Total CPU. The overall processor utilization of this project.

## License Project Viewer

Go to the **License Project Viewer** by clicking **By License Project** under the Job Info section of the Grid menu bar. Resources are shown in a cluster by project.

The information that is shown on the page is as follows:

- License Project. The name of the license project. Click a license project name to show running jobs for this license project (on the **Job Info > Details** page).
- Cluster Name. The LSF cluster to which this license project belongs.
- Total Slots. The total number of job slots that are used for this license project.
- Pending Slots. The number of job slots that are used by pending jobs for this license project.
- Running Slots. The number of job slots that are used by running jobs for this license project.
- Avg Effic. The average efficiency of this license project.
- Max Mem. The maximum memory that is used by this license project.
- Avg Mem. The average amount of memory that is used by this license project.
- Max Swap. The maximum swap space that is used by this license project.
- Avg Swap. The average swap space that is used by this license project.
- Total CPU. The overall processor utilization of this license project.

## Queue Viewer

Go to the **Queue Viewer** by clicking **By Queue** under the Job Info section of the Grid menu bar. This display is similar to the LSF **bqueues** command, with these exceptions: It includes the average and maximum run time of jobs in that queue, and the average and maximum pending time for the queues.

The information that is shown on the page is as follows:

- Queue Name. The name of the LSF queue. Click a queue name to show running jobs in this queue (on the **Job Info > Details** page).
- Cluster Name. The LSF cluster to which this queue belongs.
- Priority. The priority of the queue.
- Status Reason. The status of the queue, with further detail about the status.
- Max Slots. The maximum number of job slots that can be used by the jobs in the queue.
- Num Slots. The total number of available slots for this queue.
- Run Slots. The number of job slots that are used by running jobs in the queue.
- Pend Slots. The number of job slots that are used by pending jobs in the queue.
- Suspend Slots. The number of job slots that are used by suspended jobs in the queue.
- AVG Pend. The average number of job slots that are held by pending jobs in the queue.

- MAX Pend. The maximum number of job slots that are held by pending jobs in the queue.
- AVG Run. The average number of job slots that are held by running jobs in the queue.
- MAX Run. The maximum number of job slots that are held by running jobs in the queue.

If you select any queue, you are directed to a display of all "RUNNING" jobs within that queue.

## View Job Array Listing page

Go to the **View Job Array Listing** page by clicking **By Array** under the Job Info section of the Grid menu bar. The page shows information similar to the LSF **bjobs -A <job\_id>** command, but also includes aggregate information for the job array as a whole.

The information that is shown on the page is as follows:

- Array ID. The job array ID.
- Job Name. The name of the job.
- User ID. The identifier of the user who submitted the job array.
- Total Jobs. The total number of jobs in the job array.
- Pending Jobs. The number of jobs that remain pending in the job array.
- Running Jobs. The number of currently running jobs.
- Done Jobs. The number of jobs that are completed without error.
- Exit Jobs. The number of jobs where errors prevented the job from completing.
- Array Effic. The average CPU efficiency of jobs in the job array.
- Avg Memory. The average memory that is used by jobs in the array.
- Avg Swap. The average swap space that is used by jobs in the array.
- Total CPU Time. The total CPU time that is used by all started jobs in the job array.

If you select any array, you are directed to a display of all "ACTIVE" and "FINISHED" jobs within the job array.

## View Job Listing page

Go to the **View Job Listing** page by clicking **Details** under the Job Info section of the Grid menu bar. Filter batch job information to view only the job types you are interested in.

You can filter your view of the data by providing a resource string that conforms to the LSF **bhosts -R** command format. This displays jobs that are running on hosts that match the resource requirement. It has no comprehension of any job-specific resource requirements.

Clear the **Dynamic** check box if you do not want to immediately update page information each time you change a filter setting, and instead want to wait until you complete all filter settings and then click **Go**.

The information that is shown on the page is as follows:



- **Job ID.** The job ID that LSF assigned to the job. Click a job number to view an information page that contains details about that job (including general job information, job submission details, the job execution environment, current/last job status, and a graphical job history).
- **Job Name.** The name of the job.
- **Status.** The status of the job.
- **State Changes.** The number of times that the status of the job was changed.
- **User ID.** The LSF user who submitted the job.
- **Mem Usage.** Total resident memory usage of all processes in the job, in MB.
- **VM Size.** Total virtual memory usage of all processes in the job, in MB.
- **CPU Usage.** CPU utilization for this job.
- **CPU Effic.** The efficiency with which this job is using the CPU allocated to it, expressed as a percentage.
- **Start Time.** The time at which the job was started.
- **Pend.** The length of time in which the job was in the pending state.
- **Run.** The length of time for which the job was running.
- **SSusp.** The length of time the job was suspended by the system.

At the bottom of the **Details** page there are color-codes that indicate job efficiency thresholds, including Warning, Alarm, Flapping, and Dependencies. You can set the colors for each of these thresholds from the **Console** tab, on the **Grid Settings > Status/Events** page, along with the thresholds themselves.

## View Jobs by Application

Go to the **Batch Application** page by clicking **By Application** under the Job Info section of the Grid menu bar. You can quickly check job details by application on the page.

By default, only applications with running jobs are listed. To view all applications, click that **Include Unused Applications** check box when you filter the search.

### Note:

- LSF Application must be configured first at:  
\$LSF\_ENVDIR/lSBATCH/<clustername>/configdir/lb.applications.
- Unused application statuses are cleared at regular intervals. You can configure statuses by going to **Console > Grid Setting > Poller** tab. Set the interval under Cluster Graph Management section.

You can click **View Active Jobs by Application** action icon to view the active jobs by application.

If you want to see the running jobs by application, click the **Application name** link.

You can also click **View Application Graphs** action icon to see the job application graphs. The following are the six new graphs:

- GRID - Applications - Efficiency
- GRID - Applications - Memory Stats
- GRID - Applications - Pending Jobs
- GRID - Applications - Running Jobs



- GRID - Applications - Total CPU
- GRID - Applications - VM Stats

## View Application Profile in Job Detail page

Go to the **Job Detail** page by clicking **Details** under the Job Info section of the Grid menu bar. Filter batch job information to view only the job types you are interested in. You can filter the jobs by application by selecting **All** or the specific application through the **Apps** field. Click **Go** after you complete the filter setting.

Double-click the **JobID** of the job you want to see more details. Click **Job Detail** tab.

On the Submission Details section of the page, the Application details are displayed.

### Note:

Only users with the View Application Data permission can see the Applications Data. Permission can be set in the realm permission for each user account.

## Purge Application information

Purge old application information by setting the application purge frequency.

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings**.
3. Click the **Poller** tab.
4. On the Cluster Graph Management section, set the **Queue Host Group and Application Purge Frequency** to how long after the application is removed from the system you want to have the corresponding graphs removed.
5. Click **Save**.

## View Jobs by Group

Go to the **Batch Job Groups** page by clicking **By Group** under the Job Info section of the Grid menu bar. You can quickly check job details by group from the page.

You can click **View Active Jobs by Job Group** action icon to view jobs by job group.

If you want to view running jobs by job group, click the corresponding **Group Name** link.

### Note:

- If you filter the Group by Level, data from hidden sublevels are shown in the upper level job group.  
The maximum level on the Level drop-down list is based on the job group aggregation level setting.
- By default, unused job groups are not displayed. You can filter your search to show unused groups by clicking the **Include Unused Groups** search criteria.

You can also click **View Job Group Graphs** action icon to show more information about the group. It displays six graphical figures that represent statistics on job

efficiency, memory stats, pending jobs, running CPU, running jobs, and VM stats. These graphs provide a general image of your jobs and you control the job flexibly through:

- GRID – Job Groups - Efficiency
- GRID - Job Groups - Memory Stats
- GRID - Job Groups - Pending Jobs
- GRID - Job Groups - Running Jobs
- GRID - Job Groups - Total CPU
- GRID - Job Groups - VM Stats

## Set Job Group Aggregation Level Limit

Set the job group aggregation level to limit the Level filter options.

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings**.
3. Click the **Aggregation** tab.
4. On the Job Group Tracking section, specify the **Aggregation Level**.

Platform RTM aggregates job group information only by the specified level. If the level is set to 0, the job group data is not aggregated.

### Note:

Only users with the View Job Group Data permission can see the Group Data. Permission can be set in the realm permission for each user account.

5. Click **Save**.

## Purge Job Group information

Purge old job group information by setting the job group purge frequency.

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings**.
3. Click the **Poller** tab.
4. On the Cluster Graph Management section, set the **User, User Group, Job Group, Project, License Project Purge Frequency** to how long after the job group is not updated you want to automatically remove the corresponding graphs.
5. Click **Save**.

## View Job Group in Job Detail page

Go to the **Job Detail** page by clicking **Details** under the Job Info section of the Grid menu bar. Filter batch job information to view only the job types you are interested in. In this release, you can filter the jobs by selecting the group through the **JGroup** field. Click **Go** after you complete the filter setting.

Double-click the **JobID** of the job you want to see more details. Click **Job Detail** tab.

On the Submission Details section of the page, the Job Group details are displayed.

## View Exception Status in Job Detail page

### Important:

You must define the following parameters in the LSF `lsb.queues` file:

- `JOB_IDLE` for idle job exception handling
- `JOB_OVERRUN` for exception handling of jobs that run longer than specified run time
- `JOB_UNDERRUN` for exception handling of jobs that exits before the specified number of minutes

After you configure these parameters, you can submit jobs to the new queue which now includes the exception status definition.

Go to the **Job Detail** page by clicking **Details** under the Job Info section of the Grid menu bar. Filter batch job information to view only the job types you are interested in. Click **Go** after you complete the filter setting

Double-click the **JobID** of the job you want to see more details. Click **Job Detail** tab.

On the Current/Last Status section of the page, the Exception Status (such as idle, overrun, or underrun) details are shown. Abnormal exit shows job exit information that is predefined in LSF.

## Enable Pending Reason History Reporting

You can collect and store all pending reasons throughout a job's lifecycle. From this information, you can identify the area where jobs are pending for the longest duration.

### Note:

- By default, the pending reason collection and aggregation feature is disabled.
  - Collecting pending reason from LSF increases the polling time.
  - For large clusters with over 100K jobs per day, you must configure `CONDENSE_PENDING_REASON` parameter in LSF. Refer to the LSF documentation for more details.
1. Click the **Console** tab.
  2. Under the Configuration section of the Console menu bar, click **Grid Settings**.
  3. Click the **Poller** tab.
  4. On the Job Collection Settings section, click **Enable Pending Reason History Reporting** check box.
  5. Click **Save**.

Pending reasons duration analysis starts and a Gantt chart displays the Pending Reason History Report on the Pending Reasons tab in the job's Details page.

## Filter pending reasons on the Settings page

You can set the Pending Reason History Report to ignore specific pending reasons.

1. Click **Grid** tab.
2. Click **Settings** tab.

3. Click **Pending Reason** tab.
4. Click the check box corresponding to the reasons you want the system to ignore.  
You can also enter the reason on the Ignore Pending Reason by textbox under the General Setting (by RegExp) section.
5. Click **Save**.

**Note:**

- You can enter the reason to be ignored from General Setting (by RegExp).
- The Pending Reason and Suspend Reason sections lists the normal pending reasons (such as job slot limit reached, load information unavailable, idle time is not long enough).
- The Load Indices section lists resource-based reasons (such as r15s 15-second load averaged over the last 15 seconds, r1m 1-minute load average over the last minute, io Disk I/O rate averaged over the last minute).

The list includes all currently known RTM pending and suspend reasons. Not all LSF implemented pending and suspend reasons are included.

## View Job Pending Reason History page

Go to the **Job Pending Reason History** page by clicking **Details** under the Job Info section of the Grid menu bar. Filter batch job information to view only the job types you are interested in. Click **Go** after you complete the filter settings.

Double-click the **JobID** of the job you want to see the pending reason history details.

Click **Pending Reasons** tab. The Job Pending Reason History page displays the Pending Reason History report.

**Note:**

- The Pending Reasons tab is only visible if a job has at least one pending reason that occurred in the past.
- By default, the report is sorted by pending duration, lists the top 20 pending reasons, and applies the ignore setting of the RTM system.
  - Pending Duration: displays reasons that are sorted by the total duration that the job spent for each reason
  - Pending Reason Text: sorts by the pending reason text
  - First Start Time of Reasons: sorts by the start time each pending reason occurred
  - The Last End Time of Reasons: sorts by the end time of each pending reason occurred
- Time header (unit) automatically adapts as second/minute/hour/day by pending reason history.
- Report shows Now vertical line if the job is pending or suspended.
- Red line is displayed in the report when the job state is changed to running state. It signifies the current time of the job when it is not yet finished.

The **Pending Reasons** tab remains even if the job is finished.

## User/Group Info section

The User/Group Info section is in the Grid menu bar.

### User Viewer page

Go to the **User Viewer** page by clicking **Users** under the User/Group Info section of the Grid menu bar. The page shows job information that pertains to an LSF user:

- **User Name.** The name of the LSF user. Click a user name to show details of running jobs that are submitted by this user (on the **Job Info > Details** page).
- **Max Slots.** The maximum number of job slots that can be processed concurrently for the specified user's jobs.
- **Num Slots.** The current number of job slots that are used by the specified user's jobs.
- **Started Slots.** The number of slots that are used by user-submitted jobs and started by LSF. Started jobs can either be running, system-suspended, or user-suspended.
- **Pending Slots.** The number of job slots that are used by the user's pending jobs.
- **Running Slots.** The number of job slots that are used by the user's running jobs.
- **Effic % %.** The average CPU efficiency for jobs that are submitted by this user.
- **Sys Susp Slots.** The number of job slots that are used by the user's system-suspended jobs.
- **User Susp Slots.** The number of job slots that are used by the user's user-suspended jobs.
- **Reserve Slots.** The number of job slots that are used by the user's pending jobs.

### User Groups Viewer page

Go to the **User Groups Viewer** page by clicking **Groups** under the User/Group Info section of the Grid menu bar. Job information that pertains to an LSF user group is shown:

- **Group Name.** The name of the LSF user group. Click a user group to show details of running jobs that are submitted by this user group (on the **Job Info > Details** page).
- **Max Slots.** The maximum number of job slots that can be processed concurrently for the jobs that are submitted by the specified user group's jobs.
- **Num Slots.** The current number of job slots that are used by the specified user group's jobs.
- **Started Slots.** The number of slots that are used by user group-submitted jobs and started by LSF. Started jobs can either be running, system-suspended, or user-suspended.
- **Pending Slots.** The number of job slots that are used by the user group's pending jobs.
- **Running Slots.** The number of job slots that are used by the user group's running jobs.
- **SSusp Slots.** The number of job slots that are used by the user group's system-suspended jobs.
- **USusp Slots.** The number of job slots that are used by the user group's user-suspended jobs.
- **Reserve Slots.** The number of job slots that are used by the user group's pending jobs.

## Load Info section

The Load Info section is in the Grid menu bar.

### Host Load Viewer

Go to the **Host Load Viewer** by clicking **Host** under the Load Info section of the Grid menu bar. The page shows information that is similar to the LSF command `lsload`.

- **Host Name.** The name of the host. Click a host name to show details about jobs that are running on that host (on the **Job Info > Details** page).
- **Cluster.** The LSF cluster to which this host belongs.
- **Type.** The type of the host, as specified in the LSF configuration.
- **Model.** The model of the host, as specified in the LSF configuration.
- **Status.** The status of the host.
- **Duration In Current State.** The amount of time that the host spent in its current state.
- **RunQ 15 sec.** The exponentially averaged effective CPU run queue length of the last 15 seconds.
- **RunQ 1 min.** The exponentially averaged effective CPU run queue length over the past 1 minute.
- **RunQ 15 min.** The exponentially averaged effective CPU run queue length over the past 15 minutes.
- **CPU %.** The current CPU utilization rate.
- **Page Rate.** The memory paging rate exponentially averaged over the last minute, in pages per second.
- **I/O Rate.** The disk I/O rate exponentially averaged over the last minute, in kilobytes per second.
- **Cur Logins.** The number of current login users.
- **Idle Time.** The idle time of the host, in minutes.
- **Temp Avail.** The amount of free space in /tmp (G=gigabyte, M=megabyte).
- **Swap Avail.** The amount of swap space available (G=gigabyte, M=megabyte).
- **Mem Avail.** The amount of physical memory available (G=gigabyte, M=megabyte).

### Group Load Viewer

Go to the **Group Load Viewer** by clicking **Host Group** under the Load Info section of the Grid menu bar. The page shows host performance information that is aggregated to the level of LSF host group:

- **Group Name.** The name of the host group. Click a group name to go to the **Host Load** page and view information similar to running the LSF command `lsload`.
- **Cluster.** The cluster to which the host group belongs.
- **Status.** The status of the host group.
- **AvgRq 15 sec.** The exponentially averaged effective CPU run queue length for hosts within the group, for the past 15 seconds.
- **AvgRq 1 min.** The exponentially averaged effective CPU run queue length for hosts within the group, for the past 1 minute.
- **AvgRq 15 min.** The exponentially averaged effective CPU run queue length for hosts within the group, for the past 15 minutes.

- Avg CPU %. The average CPU utilization for hosts within the group.
- Avg Page Rate. The average memory paging rate exponentially averaged over the last minute for hosts within the group, in pages per second.
- Avg I/O Rate. The average disk I/O rate exponentially averaged over the last minute for hosts within the group, in kilobytes per second.
- Total Logins. The total number of current login users for hosts within the group.
- Avg Idle Time. The idle time of the host, in minutes.
- Avg Temp Avail. The average amount of free space in /tmp (G=gigabyte, M=megabyte).
- Avg Swap Avail. The average amount of swap space available for hosts in the group (G=gigabyte, M=megabyte).
- Avg Mem Avail. The average maximum amount of physical memory available for user processes on the host (G=gigabyte, M=megabyte).

You can view Job Statistics graphs by clicking the corresponding **View Host Group Graphs** action icon of the Host Group. These graphs present images and statistics to help you control jobs flexibly.

- GRID - Host Group - CPU Utilization  
Presents statistics on CPU utilization when the host is the server and its status is not Unavail, Unreach, Unlicensed, or Closed-Admin.
- GRID - Host Group - Available Memory  
Presents a general image of available memory by host group level for more flexibility in controlling jobs.
- GRID - Host Group - Host Details  
Presents a general image of memory usage by host group level.
- GRID - Host Group - Memory Stats  
Presents statistics on average or maximum memory when job status is RUNNING, USUSP, and SSUSP.

## Host Info section

The Host Info section is in the Grid menu bar.

### Closed Admin Host Viewer

Go to the **Closed Admin Host Viewer** by clicking **Closed** under the Host Info section of the Grid menu bar. The page shows information about closed admin hosts such as Host Name, Cluster, Host Control Message, CPU Pct, RunQ 1m, Mem Usage, Page Usage, and Page Rate details. Select the host then choose **Open** action to view more details.

### Queue Host Distribution Matrix Viewer

Go to the **Queue Host Distribution Matrix Viewer** by clicking **QDistrib** under the Host Info section of the Grid menu bar.

#### Note:

- The page shows the distribution of hosts across different queues in a matrix format.  
The first user who accesses the page might see a blank page. The matrix must be rebuilt to see the details.
- Click the **Rebuild** button to rebuild the matrix from the current configurations.



This table is rebuilt on demand. The bottom of the table shows the **Table Last Rebuilt** details.

- The page is controlled by realm permissions and is disabled by default.
- Filters for **Cluster**, **Model**, **Group**, **Type**, **Refresh**, and **Records** are available on the page. Click **Go** to search.
- Select **HG: Roll-Up** from the Group filter list to see a summarized report of used and total slot counts by Host Group. You can also enter values to search for groups of queues.

The page shows corresponding Action icons (such as View Host Graphs and SSH - Remote to Host) and details for Host, Cluster, Use Type, Max Slots, Model, Type, Memory, Running, Status, Chkpnt/Rerun/Queue, Idle, License, Night, Normal, Owners, Priority, and Short columns in the table matrix.

## Host Status Viewer

Go to the **Host Status Viewer** by clicking **Servers** under the Host Info section of the Grid menu bar. The page shows information identical to the LSF **lshosts** command.

- Host Name. The name of the host.
- Cluster. The LSF cluster to which this host belongs.
- Type. The type of the host, as defined in the LSF configuration.
- Model. The model of the host, as defined in the LSF configuration.
- CPU Factor. The CPU factor of the host, as defined in the LSF configuration.
- Max CPUs. The number of processors on this host.
- Max Memory. The maximum amount of physical memory available for user processes (G=gigabytes, M=megabytes).
- Max Swap. The total available swap space (G=gigabytes, M=megabytes).
- Max Temp. The maximum available space in /tmp (G=gigabytes, M=megabytes).
- Total Disks. The number of local disk drives directly attached to the host.
- Resources. The Boolean resources that are defined for this host, denoted by resource names, and the values of external numeric and string static resources.

## Client Viewer

Go to the **Client Viewer** by clicking **Clients** under the Host Info section of the Grid menu bar. The page shows information similar to the LSF command **lshosts** except that it displays only LSF clients. In addition to showing current clients that are registered with LSF, it shows all prior clients that performed operations on the LSF Cluster.

- Client Name. The host name of the client.
- Cluster. The cluster to which this client belongs.
- Client Type. The type of LSF client. Clients are set as Fixed Client or Floating Client.
- First Seen. The date and time at which this client was first seen in this cluster.
- Last Seen. The date and time at which this client was last seen in this cluster.

The page is helpful in situations where you are using either submission-only clients or floating client configurations.



## Host Group Viewer

Go to the **Host Group Viewer** by clicking **Groups** under the Host Info section of the Grid menu bar. The page shows basic host information for each host within a host group.

- Group Name. The name of the host group.
- Cluster. The LSF cluster to which this host group belongs.
- Host Name. The host name of a host that belongs to the host group.
- Host Type. The type of the host.
- Host Model. The model of the host.
- CPU Factor. The CPU factor of the host.
- Max CPUs. The number of processors on the host.
- Max Mem. The maximum amount of physical memory available for user processes on the host (G=gigabytes, M=megabytes).
- Max Swap. The total available swap space on the host (G=gigabytes, M=megabytes).
- Max Tmp. The maximum amount of space in /tmp for the host (G=gigabytes, M=megabytes).

## Reports section

The Reports section is in the Grid menu.

### Shared Resources

Go to the **Shared Resources** page by clicking **Shared Resources** under the **Reports** section of the **Grid** menu.

The report shows all the shared resources defined in the LSF cluster. For numeric shared resources it reports on the total number, amount currently available and reserved resources.

### Resource Pools

Go to the **Resource Pools** page by clicking **Resource Pools** under the **Reports** section of the **Grid** menu.

It is the reporting on the guaranteed SLA's defined on the LSF cluster.

### Daily Statistics

Go to the **Job History Daily Statistics** page by clicking **Daily Statistics** under the **Reports** section of the **Grid** menu. The page shows daily statistics for your clusters.

- Cluster Name. The name of the LSF cluster.
- User Name. The name of the cluster user.
- Queue Name. The name of the queue.
- Project Name. The name of the project.
- Exec Host. The name of the execution host.
- Result. The result of the last job that is submitted to the execution host.
- Total Jobs. The total number of jobs submitted.
- Total Slots. The total number of slots available.

- Avg W-Time. Average job wait time.
- Total W-Time. Total job wait time.
- System Time. The amount of system that is used by submitted jobs.
- User Time. The amount of user time that is used by submitted jobs.
- Start Date. The start of the reporting period.
- End Date. The end of the reporting period.

## Parameters page

Go to the **Parameters** page by clicking **Parameters** under the Reports section of the Grid menu. The page shows defined configuration parameters for your cluster as defined in the `lsb.params` file.

- Name. Name of the parameter that is used in this cluster.
- Cluster. Name of the cluster to which the parameter applies.
- Description. Brief description of the parameter.
- Value. Configured parameter value.

## Settings tab (from the Grid page)

Go to the **Grid > Grid User Settings** tab by clicking the **Grid** tab, then by clicking the **Settings** tab on the right side of the tabbed interface. There is a subtab for each category of grid display settings that you can change.

- **General:** Define the default page to display, and the default LSF cluster for which you want to filter data.
- **Visual:** Control a number of visual settings such as the size of the summary icons and the frequency of refresh intervals.

The page contains the following fields:

- Screen Refresh Interval: The interval that is used by RTM to refresh a page you are currently viewing.
- Exclusion Filter Status: Determines whether the exclusion filter is on or off by default. If this option is checked, only hosts that match the exclusion filter settings in the field are shown in the host dashboard.
- Exclude Host States: Select one or more of the exclusion states to eliminate hosts of that state from displaying in the dashboard. Doing so leaves only the hosts with states you care about in the dashboard. For it to work, you must enable (check) Exclusion Filter Status.

To select multiple states, press and hold the <CTRL> key while you choose host states from the list.

- **Timespans:** Control how your grid summary icons are displayed. Each user can display thresholds that they believe represent how their grid hosts are behaving. Among the fields that are displayed, take note of the following specific field:
  - Default Grid View Timespan: When you view jobs in the RTM Interface, indicates which default Timespan is in effect. This setting applies only when you view the following status names: “All”, “Done”, and “Exit”.
- **Clusters:** Customize your view of the Cluster Dashboard (**Dashboards > Cluster** page). Choose to show or hide various information columns.
- **Hosts:** Show or hide certain fields on the **Host Details** page.
- **Host Groups:** Show or hide certain fields on the **Host Details** page.
- **Queues:** Show or hide certain fields on the **Job Info > By Queue** page.

- **Jobs:** Show or hide certain fields on the **Job Info > Details** page. By default, not all column details are shown in the Details page. As there is a large quantity of data that is collected on LSF jobs, you might want to show only the information that you are most interested in.

List of Job Details column includes: Cluster ID, Cluster Name, Job Name, Project Name, License Project, Application, and Job Group.

- **Job Export:** Show or hide certain fields on the **Job Info > By Host** and **Job Info > By Host Group** pages. As there is a large quantity of data that is collected on LSF jobs, you might want to show only the information that you are most interested in.
- **Arrays:** Customize your view of the **Job Info > By Array** page. Choose those fields that are most relevant for your environment.
- **Job Graphs:** Customize the display of your **Job Info > Details** graph according to your preference.  
You might also customize the colors that RTM uses when you are constructing job graphs. The lists in the **Graph Colors** section are color-coded to facilitate your choice.
- **Pending Reason:** Customize the display of the **Job Pending Reason History** page. Choose to show or hide RTM pending and suspended job reasons. The list shows only pending and suspend reasons that were processed by RTM.

---

## View time series graphs for clusters and hosts

Time series graphs present cluster status and details that depend on time. Time is the X-axis in these line graphs.

The following is a list of actions you can take to view time series graphs of your cluster status and details:

- View graphs by using the Tree, List, or Preview view.  
Click the **Graphs** tab, then one of the icon tabs on the right side of the tabbed menu interface to view time series graphs in the **Graphs** page. For more information, see “Viewing modes.”
- View graphs by using the action icons.  
When you are viewing rows of data that have action icons, you can view time series graphs for the specific row (such as a specific cluster, host, or queue) by clicking the **View host job detail** or **View graphs action** icon.
- Configure the graph display settings.  
Click the **Graphs** tab, then the **Settings** tab on the right side of the tabbed menu interface. For more information, see “Settings tab (from the Graph page)” on page 40.

## Viewing modes

When in the **Graphs** page, you can view graphs by using the Tree, List, or Preview views.

### Tree view

Click the **Graphs** tab, then the **Tree view** tab on the right side of the tabbed menu interface to access the graphs in Tree Mode. Use this viewing mode to access all graphs as organized by device and cluster in the tree. This tree is customizable, as are all graphs, from the **Console** tab. Devices can be added to the tree and modified from the **Console** tab.

What you can do:

- Click a root-level tree name to view summary information about all devices within that section (view thumbnails).
- Click a tree section to view host-specific information in graph form.
- Beside a graph, click the magnifying glass to view more detailed information that is broken down in different ways.

If you zoom into a graph, you are presented with detailed information.

- Beside each graph, you can click a wrench icon to display helpful debugging information that is collected by an RRD tool.

## List view

Click the **Graphs** tab, then the **List view** tab on the right side of the tabbed menu interface to access the graphs in List Mode. Select one or more cluster names from the list, then click **View**. The selected graphs display.

## Preview

Click the **Graphs** tab, then the **Preview** tab on the right side of the tabbed menu interface to access the graphs in Preview Mode. From the preview page, you can filter by host to limit the number of graphs displayed.

## Settings tab (from the Graph page)

Go to the **Graphs > Settings** page by clicking the **Graph** tab, then by clicking the **Settings** tab on the right side of the tabbed interface. You can configure the appearance of your graphs and default page settings, and the page consists of the following sections:

- **General:** Configure general graph options and display formats.
- **Graph thumbnails:** Configure the size of the thumbnails that are used to represent your graphs.
- **Tree View Mode:** Configure the tree defaults and display when in tree view mode.
- **Preview Mode:** Configure the preview mode display.
- **List View Mode:** Configure the list view mode display.
- **Graph Fonts:** Choose whether to use your own custom fonts and font sizes or the system defaults.

---

## View non-time series charts for clusters and hosts

### About this task

Non-time series charts present cluster and host statistics that do not depend on time. Because the statistics are not time-sensitive, Platform RTM can present this data in graphically rich formats, such as pie charts and bar graphs.

The following describes the steps that you must take to view non-time series charts for clusters and hosts:

### Procedure

1. Click the **Grid** tab.
2. Under the Dashboards section of the Grid menu bar, click **Statistical**.

3. In the **Cluster Statistics** selection filter, choose the types of data that you would like to view as charts.

If there is a **Selectable Reports** dialog next to the charts (such as the **User Stats** chart in the previous example), select different types of reports to view different data by clicking the name of the report.



---

## Chapter 3. Working with application license servers

You can use Platform RTM to configure and monitor application license servers.

---

### Relationship between Platform RTM and license servers

Platform RTM is used to monitor and graph LSF resources (including networks, disks, and applications) in a cluster. In graph or table formats, Platform RTM displays resource-related information such as the number of jobs that are submitted, the details of individual jobs (like load average, CPU usage, job owner), or the hosts on which the jobs ran.

FLEXlm is a third-party license manager that is monitored by Platform RTM. FLEXlm allows licenses to remain on the network instead of a specific host; any registered host in the cluster can then access and use a specific licensed application, increasing usage efficiency.

---

### License Management section

The **License Management** section is in the Console menu bar.

#### License Pollers page

Go to the **RTM License Pollers** page by clicking **License Pollers** under the **License Management** section of the Console menu bar. The page shows information about the pollers that collect data from the license servers.

- **Add.** Add a license poller for Platform RTM. Click to open the **RTM License Poller Edit [new]** page and specify the properties of the license poller.
- **License Poller Name.** The defined name for the license poller. Click a name to open the **RTM License Poller Edit** page and edit license poller properties.
- **License Poller ID.** The ID assigned to the poller.
- **License Poller Hostname.** The host name of the server that contains the remote poller.
- **License Poller Path.** The file path to the license poller binary file.
- **Department.** The specified department responsible for this license server.
- **License Poller Type.** The type of license server (for example, FLEXlm).
- **Choose an action.** Select one or more check boxes for the servers on which to act (for example, **Delete**), then select an action and click **Go**.

#### License Services page

Go to the **License Services** page by clicking **License Services** under the **License Management** section of the Console menu bar. The page shows information about the license services.

- **Add.** Add a license service for Platform RTM to monitor. Click to open the **License Services Edit [new]** page and specify the properties of the license service.

The **License Services Edit [new]** page can be used to set the Daemon Type, FLEXlm General Information (such as License Service Name, License Poller, Service Port@Server Notation, and License Service Timezone), FLEXlm

Connection Settings (such as Monitor License Checkouts, Poller Interval, Connection Timeout, and Connection Retries), and FLEXlm Support Information (such as Service Vendor, Location, Region, Department, License Type, and Support Contacts).

Based on the settings, the filter criteria becomes available on the License Services pages, License Inventory, and Feature Usage pages. You can quickly search for license services by using filter criteria.

- License Service Name. The defined name for the license service. Click a name to open the **License Service Edit** page and edit server properties, connection settings, and support information.
- License Service ID. The ID assigned to the license service.
- License Poller Name. The name of the poller that is associated with this service.
- Poller Interval. The license poller interval.
- Vendor. The vendor for the license service.
- Department. The Department responsible for the license service.
- Type. This License Type column shows the license service purpose.
- Location. The physical location of the license service.
- Region. The region for the license service.
- Collect Status. The status of the license service. (Status can be Down, Recovering, Up, and Unknown.)
- Current Time. The time that is taken to poll this service.
- Average Time. The average time that is taken to poll this service.
- Max Time. The maximum time that is taken to poll this service.
- Availability. The availability of the license service that is based on the percentage of successful polling, reported as a percentage.
- Last Failed. The last time that the license poller failed to poll the license service. If there was no failure, N/A is shown.
- Last Updated. The last time that the license poller successfully polled the license service. If the license poller never successfully polled the license service, N/A is shown.
- Choose an action. Select one or more check boxes for the service on which to act (for example, **Delete**, **Disable**, **Enable**, or **Clear Stats**), then select an action and click **Go**.

---

## Configure license pollers

License pollers allow Platform RTM to monitor license servers. The following describes the actions that you must take, in chronological order, to configure license pollers.

### Set up remote license pollers

Optional. Platform RTM is preconfigured with local license pollers to collect information from license servers. If you want to use remote license pollers, you must set up and configure the remote license pollers to work with Platform RTM. For more details on how to set up and configure remote license pollers on Platform RTM, refer to the `INSTALL.REMOTE_LICPOLLER` guide.

### Add or edit license pollers to the Platform RTM Console

Platform RTM uses license pollers to collect information from license servers.



## Procedure

1. Click the **Console** tab.
2. Under the License Management section of the Console menu bar, click **License Pollers**.

The **Pollers Console** page displays.

3. Add or edit a license poller.
  - To add a license server, click **Add**.
  - To edit an existing license server, click the name of the license server that you want to edit.

The **Platform RTM License Poller Edit** page displays.

4. Specify (or update) the required fields that describe your license poller.
5. Click **Create** (or **Save**) to save the settings for your license poller.

---

## Add or edit license servers for Platform RTM to monitor

### About this task

You must add a license server to the Platform RTM host for Platform RTM to monitor the license server and create the necessary license graphs. Platform RTM supports FLEXlm license server.

## Procedure

1. Click the **Console** tab.
2. Under the License Management section of the Console menu bar, click **License Services**.

The **License Servers Console** page displays.

3. Add or edit a license server.
  - To add a license server, click **Add**.
  - To edit an existing license server, click the name of the license server that you want to edit.

The **License Server Edit** page displays.

4. Specify (or update) the required fields that describe your license server.
  - For the **License Poller** field, select the appropriate poller for your version of the LSF cluster.
  - For the **Daemon Type** field, select the type of license server in the LSF cluster (FLEXlm).
5. Click **Create** (or **Save**) to save the settings for your license server.

### What to do next

- If you edited a license server that was already in the Platform RTM Console, you do not have to do anything else.
- If you added a license server to the Platform RTM host, create graphs for your Cacti host device and its associated license server as described in the next section.

---

## Monitor license servers

### About this task

The following is a list of actions you can take to monitor license servers by using the Platform RTM Console.

### Procedure

- “View the status of the license services”
- “View reports on license servers and projects” on page 47

## View the status of the license services

Use the License Dashboard to view the status of all the license services that Platform RTM is monitoring.

### Procedure

1. Click the **License** tab.
2. In the Dashboards section of the License menu bar, click **License Services**.

The **License Services Dashboard** displays the name and details of the license services.

The license dashboard shows the following information:

- **Actions.** Icons that you can click include the following actions: View inventory, View charts, and View license checkouts.
  - **Service Name.** The defined name for the license service. Click **Usage Charts** from the navigation tree to see the **Usage** page and view information about licenses in the license service.
  - **Vendor.** The vendor for the license service.
  - **Department.** The department responsible for the license service.
  - **Type.** The license service purpose.
  - **Location.** The specified physical location of the license service.
  - **Collect Status.** The status of the license service. (Status can be Down, Recovering, Up, and Unknown.)
  - **Quorum Status.** The current quorum. (Status can be Down, Recovering, Up, and Unknown.)
  - **Current Time.** (Current response time) - (**rtm run lmstat** command's response time)
  - **Average Time.** (Average response time) - (**rtm run lmstat** command's response time)
  - **Max Time.** (Maximum response time) - (**rtm run lmstat** command's response time)
  - **Total Features.** The total number of features that the license service supports.
  - **In Use Features.** The number of features on the license service that are currently in use.
  - **In Use Licenses.** The number of licenses on the license service that are currently in use.
  - **Availability.** The availability of the license service that is based on the percentage of successful polling, reported as a percentage.
  - **Last Failed.** The date when the license service last failed.
3. To view more details, you can click the corresponding action icon or click the link items from the navigation tree.

## View reports on license servers and projects

The **FLEXlm Reports** section of the License menu bar has links to the relevant license reports.

### Procedure

1. Click the **License** tab.
2. Click the links in the FLEXlm Reports section of the License menu bar to view the license reports.
  - The **License Inventory** link shows license details for license server. For more information, see “License Inventory.”
  - The **Feature Usage** link shows the licenses in use for each license feature. For more information, see “Feature Usage.”
  - The **Feature Checkouts** link shows information about licenses that are currently checked out. For more information, see “Feature Checkouts” on page 48.
  - The **Usage Charts** link shows the non-time series charts for your licenses. For more information, see “Viewing non-time series charts for license servers” on page 49.

### FLEXlm Reports section

The FLEXlm Reports section is in the License menu bar.

#### License Inventory:

Go to the **License File Details** page by clicking **License Inventory** under the FLEXlm Reports section of the License menu bar. The page shows license details for license server.

- **Actions.** Icons that you can click include the following actions: View inventory, View charts, and View license checkouts.
- **Feature Name.** The license name.
- **Vendor.** The license vendor.
- **Service Name.** The name of the license server.
- **Location.** The specified physical location of the license service.
- **Department.** The department responsible for the license service.
- **Type.** The license service purpose.
- **Daemon.** The vendor daemon.
- **Version.** The version number of this licensed product.
- **Number Expiring.** The number of expiring licenses for this product.
- **Expiration Date.** The date that these product licenses expire.
- **Status.** The status of this product license.

#### Feature Usage:

Go to the **License Usage** page by clicking **Feature Usage** under the Reports section of the Grid menu bar. The page shows license usage, provided by the FLEXnet servers.

- **Actions.** Icons that you can click include the following actions: View feature graphs, View charts, and View license checkouts.
- **Feature Name.** The feature license name.
- **Vendor Name.** The name of the vendor.

- Location. The specified physical location of the license service.
- Department. The department responsible for the license service.
- Type. The license service purpose.
- Daemon. The vendor daemon.
- Daemon Version. The version number of the vendor daemon.
- Daemon Status. The vendor daemon status (status can be Up or Down).
- Total. The total number of features that the license service supports.
- In Use. The number of features that are checked out from the license service.

#### Feature Checkouts:

Go to the **License Checkouts** page by clicking **Feature Checkouts** under the Reports section of the Grid menu bar. The page shows information about licenses that are currently checked out licenses.

- Feature Name. The feature license name.
- Vendor Name. The name of the license vendor.
- Server Name. The name of the license server.
- Version. The version of the license.
- User Name. The user that checked out the license.
- Host Name. The name of the LSF server that checked out the licenses.
- Status. The status of the license checkout.
- InUse. The number of licenses that are checked out.
- Date. The time of the last change in status.

## Create graphs to monitor your license server

### Procedure

1. Click the **Console** tab.
2. Under the Create section of the Console menu bar, click **New Graphs**.  
The **Create New Graphs** page displays.
3. Select your license server in the **Host** field.
4. To filter the available types of graphs that you want to create, select a category of graphs in the **Graph Types** field.
5. In the **Create** field at the bottom of the Graph Templates column, select the type of graph you want to create for the license server.

### What to do next

Edit these graphs by clicking **Graph Management** in the Management section of the Console menu bar and clicking the name of the host and graph.

#### Note:

For more information about Data Query, see the Cacti documentation at [www.cacti.net/downloads/docs/html/](http://www.cacti.net/downloads/docs/html/).

---

## View time series graphs for license servers

Time series license graphs present license status and details that depend on time. Time is the X-axis in these line graphs.

## About this task

View time series license graphs as you would for any other time series graphs that you use to monitor your clusters. For more information, see “View time series graphs for clusters and hosts” on page 39.

**Note:** If you are in Tree View or Preview View, remember to click **List View** on the upper right corner of the page before you run the following procedures.

### Procedure

1. Click the **Graphs** tab.
2. Optional. Filter specific hosts, graph templates, and graph names.
  - Filter specific hosts by selecting a host in the **Host** field.
  - Filter graph templates by selecting a graph template in the **Template** field.
  - Filter graph names by specifying text in the **Search** field and clicking **Go**.
3. Select the check boxes beside any graph that you want to view.
4. Click **View** to view all the graphs that you selected.

---

## Viewing non-time series charts for license servers

Non-time series charts present license statistics that do not depend on time. Because the statistics are not time-sensitive, IBM Platform RTM can present this data in graphically rich formats, such as pie charts and bar graphs.

## About this task

The following task describes the steps to take to view non-time series charts for clusters and hosts.

### Procedure

1. Click the **License** tab.
2. Under the FLEXlm Reports section of the License menu bar, click **Usage Reports**.
3. Use the **License Filters** selection filter to filter information that you want the charts to display and click **Go**.



---

## Chapter 4. Setting up your email notifications for alerts

You can use Platform RTM for setting up email notifications for threshold, grid, and syslog alerts

---

### Working with thresholds and alerts

You can use Platform RTM to configure, view, and acknowledge thresholds and alerts.

#### Configure thresholds and alerts

A threshold triggers an alert if your clusters, hosts, queues, or jobs meet the conditions of the threshold.

The following is a list of the actions you can take to configure thresholds and alerts:

- “Create a threshold to trigger alerts”
- “Modify threshold settings” on page 52
- “Delete thresholds” on page 53

#### Create a threshold to trigger alerts

##### About this task

You can create thresholds through the following sources:

##### Procedure

- “Create a threshold from a graph template”
- “Create a threshold from a host” on page 52

##### Create a threshold from a graph template:

Create a threshold by using a graph template as the source.

##### Procedure

1. Click the **Console** tab.
2. Under the Management section of the Console menu bar, click **Thresholds**.
3. Click **Add** on the upper right side of the Threshold Management page.
4. In the **Source** field, select **Graph Template**.
5. Select the appropriate host name and graph template for the new threshold and click **Create**.
6. Specify the threshold values for which you want to trigger an alert and click **Create**.
7. In the Data Source Item page, make any further changes to your threshold configuration.

If threshold conditions are met, the Event Triggering sections can configure threshold event triggering, which specifies actions (commands, shell scripts, or host-level actions) to take.

- High threshold: If the threshold is breached because the data source exceeds this value, the threshold triggers the specified action.

- Low threshold: If the threshold is breached because the data source drops below this value, the threshold triggers the specified action.
- Norm threshold: If the threshold is breached, then returns to normal, the threshold triggers the specified action.

8. Click **Save** to create a threshold.

### Create a threshold from a host:

Create a threshold by using a host as the source.

#### Procedure

1. Click the **Console** tab.
2. Under the Management section of the Console menu bar, click **Thresholds**.
3. Click **Add** on the upper right side of the Clusters page.
4. In the **Source** field, select **Host**.
5. Select the appropriate host name for the new threshold.
6. In the **Graph** field, specify the graph for which you want your threshold to monitor.

The Data Source field displays, followed by the graph that you specified.

7. In the **Data Source** field, specify the data source item that you want your threshold to monitor and click **Create**.

If threshold conditions are met, the Event Triggering sections can configure threshold event triggering, which specifies actions (commands, shell scripts, or host-level actions) to take.

- High threshold: If the threshold is breached because the data source exceeds this value, the threshold triggers the specified action.
- Low threshold: If the threshold is breached because the data source drops below this value, the threshold triggers the specified action.
- Norm threshold: If the threshold is breached, then returns to normal, the threshold triggers the specified action.

8. In the Data Source Item page, make any further changes to your threshold configuration.
9. Click **Save** to create a threshold.

### Modify threshold settings

You can change the configuration of thresholds.

#### Procedure

1. Click the **Console** tab.
2. Under the Management section of the Console menu bar, click **Thresholds**.
3. Click the name of the threshold that you want to modify.
4. In the Data Source Item page, you can change your threshold configuration.

If threshold conditions are met, the Event Triggering sections can configure threshold event triggering, which specifies actions (commands, shell scripts, or host-level actions) to take.

- High threshold: If the threshold is breached because the data source exceeds this value, the threshold triggers the specified action.
- Low threshold: If the threshold is breached because the data source drops below this value, the threshold triggers the specified action.



- Norm threshold: If the threshold is breached, then returns to normal, the threshold triggers the specified action.
5. Click **Save** to apply your changes to the threshold configuration.

## Delete thresholds

### About this task

Delete thresholds when you no longer need the alerts that they trigger.

### Procedure

1. Click the **Console** tab.
2. Under the Management section of the Console menu bar, click **Thresholds**.
3. Click the check box at the right side of each threshold that you want to delete.
4. In the **Choose an action** field, select **Delete** and click **Go**.

## View threshold and alert information

Click the **Thold** tab to view threshold and alert information.

For more information about configuring thresholds, see “Configure thresholds and alerts” on page 51.

The following is a list of actions you can take to view threshold and alert information.

- Review the configured thresholds in your cluster.  
Click the **Thold** tab, then the **Thresholds** subtab. For more information, see “Thresholds subtab.”
- Review the status of the hosts that are being monitored by a threshold.  
Click the **Thold** tab, then the **Host Status** subtab. For more information, see “Host Status subtab” on page 54.

### Thresholds subtab

Go to the **Thresholds** page by clicking the **Thold** tab, then the **Thresholds** subtab. The page shows the configured thresholds in your cluster.

The page contains the following fields:

- Name. The name of the cluster and the threshold. Click the name of a threshold with triggered alerts to see a list of the hosts and the specific data source values that triggered the alert.
- ID. The ID assigned to the threshold.
- Type. The type of threshold (for example, High/Low, Baseline, or Time Based)
- High. The high threshold boundary value. If the current value of the monitored data source is greater than this boundary, the threshold triggers an alert.
- Low. The low threshold boundary value. If the current value of the monitored data source is lower than this boundary, the threshold triggers an alert.
- Current. The current value of the monitored data source.
- Enabled. Indicates whether this threshold is active.

At the bottom of the **Thresholds** page are color codes that indicate threshold conditions, including Alarm, Warning, Notice, Ok, and Disabled.

## Host Status subtab

Go to the **Host Status** page by clicking the **Thold** tab, then the **Host Status** subtab. The page shows the status of the hosts that being monitored by a threshold.

The page contains the following fields:

- Description. A description of the host. This description is the same as the host name for automatically added LSF hosts.
- ID. The host or device ID.
- Graphs. The number of graphs for the host.
- Data Sources. The number of data sources for the host.
- Status. The status of the host.
- Event Count. The number of threshold-triggered alerts.
- Hostname. The name of the host.
- Current (ms). The current host ping time, in milliseconds.
- Average (ms). The average host ping time, in milliseconds.
- Availability. The percentage of time that the host is available.

## Acknowledge triggered alerts for thresholds

After a threshold triggers an alert, you can acknowledge the alert to prevent future email and syslog notifications.

- “Acknowledge alerts for a single threshold”
- “Acknowledge alerts for multiple thresholds”

### Acknowledge alerts for a single threshold

#### Procedure

1. Click the **Thold** tab.

If there are several thresholds, you can use the Threshold Status menu bar to filter the threshold view.

2. In the **Actions** column, click the acknowledge icon next to the name of the threshold with the triggered alerts.

The acknowledge icon changes into the reset acknowledge icon.

### What to do next

You can click the reset acknowledge icon next to the threshold to allow the threshold to resend future email and syslog notifications with each triggered alert for the threshold.

### Acknowledge alerts for multiple thresholds

#### Procedure

1. Click the **Console** tab.

If there are several thresholds, you can use the Threshold Status menu bar to filter the threshold view.

2. Under the Management section of the Console menu bar, click **Thresholds**.
3. Click the check box at the right side of each threshold with triggered alerts that you want to acknowledge.
4. In the **Choose an action** field, select **Acknowledge** and click **Go**.
5. In the Acknowledge Message window, specify an acknowledgement reason message (or leave blank for no message) and click **Yes** to acknowledge the triggered alerts for all thresholds.

This message is recorded in the `cacti.log`, threshold log database table, and syslog files.

### What to do next

You can repeat the preceding steps, but select **Reset Acknowledgement** in the **Choose an action** field to allow the thresholds to resend future email and system log notifications with each triggered alert for the threshold.

---

## Working with grid alerts and templates

Configuring alerts is an advanced feature. You must have a good knowledge of the RTM database and structure to be able to use this advanced RTM feature for generating generic alerts. You can create alerts that are based on grid data.

To view Alerts, navigate to **Console > Management > Alerts** or **Grid > Dashboards > Alerts**.

### Alert Tabs

You can view the alert details by clicking an Alert name on **Console** or by clicking the **Edit Alert** in the **Grid** page. Current items that satisfy the conditions of your expressions or scripts are shown on the page. If you are creating a template, then configure the following five tabs:

- **General:** Define basic information such as name, severity, and type. After you save the template, the **Data Source** tab is displayed.
- **Data Source:** Define how data is retrieved to evaluate if there is an Alert. The data sources are of three types:
  - **Expression**  
A GUI builder can create expressions that are to be evaluated for threshold value.
  - **SQL Query**  
You can create your own SQL query statements.
  - **Script**  
You can create custom scripts that return threshold values. Your script can have complex queries.
- **Actions:** Define email and syslog notifications.
- **Layout::** Define which columns are to be included in both the Alert Log and alert emails.
- **Breached Items:** View the current breached alerts.

The generic alert can pull data from any database table and create threshold values for it.

**Restriction:** Each alert can operate only on one database table. The Expression builder does not allow creating expressions that span more than one table.

You can build a generic alert by associating an expression with it and by determining its threshold (Hi/Low values). You can also associate the alerts with your own custom scripts.

**Important:** Grid Alerts are controlled by realm permission. To set permissions, click **User Management** under the Utilities section of the Console menu. In the **Realm Permissions** tab, you can select or clear the permission types to enable or disable user access accordingly:

RTM > **Configure Grid Alerts** or RTM > **View Alerts**

## Creating an Alert Template

### About this task

#### Procedure

1. Go to **Alert Templates** in the **Template** section.
2. Click **Add** on the upper right side.
3. Under the **General** tab, enter the required values and save the template to see the **Data Source** tab.
4. Go to the **Data Source** tab and complete the required fields and **Save**.

You can choose one of the following expression types from the **Data Type** menu and click Save to see other options:

- **Table** - Grid table names are listed from the database. Choose a table to base the alarm.
  - a. Create **Data Source Metrics** and **Custom Data Input** Items if you want to use them for building expressions in Data Source Items.
  - b. Click **Add** in **Data Source Items** to build your expression.
  - c. Click **Save**.

The created expression displays on the **Expression Items** list. You can delete or edit the selected expression item. After you create an expression, you cannot change the database table that is associated with it.

- **SQL Query** - You can use to define your own database query.

If you are using SQL Query type Data Sources, then the Data Source returns rows of information about exceptional information in the cluster. When you define an Alert, you can select a column to aggregate within the table. For example, if you want an Alert to trigger when there are more than 1000 CPU allocated to jobs that are running over two days. Then, return a table of subset jobs from the grid\_jobs table that ran for more than two days. In that table, include the column num\_cpus, then aggregate it as a SUM, and then either use a High/Low or Time Based threshold with a breach value of 1000.

- Click **Check Syntax** to verify expression or click **Get SQL** to check your SQL query.

Either OK or Bad Syntax result is displayed. Refer to the Cacti log for the actual SQL error.

- **Script** - You can create alarms with your own scripts. These scripts are useful when you need an alarm that requires information from more than one database table or some additional process or aggregation.

If you are using Script type Data Source, then the Data Source returns two objects, a breach value and data to be processed by the Alert. So, for this type of Data Source, there is no aggregation available, the script must aggregate internally. With a Script Data Source, the data that returned can be:

- CSV with a Header,
- CSV without a Header,
- HTML, or

- Text

The returned data is included in the Alert log, and any email sent to Administrators or users.

For this method, you must provide these two items when you are configuring the Alarm Settings:

- **Threshold Script** - The script that must be run to get the current value of the alarm. You must specify the path to this threshold script. Maximum length of the script's path, name, and arguments must not exceed 255 characters.
- **Breached Item Data Script** - The script that must be run to get the breached items. The output can be in plain text or HTML. RTM automatically includes the output into the email alert or the Alarm Details page.

The maximum length of the script's path, name, and arguments must not exceed 255 characters.

- Go to the **Actions** tab and configure email notification and threshold settings.
  - If you are using expressions, the details for the **Email Subject** and **Email Body** contents are created by RTM.
  - The **Notify Accounts** can append or remove items from the two multi-select boxes while it specifies the accounts to be notified when the alert is breached. In **Extra Alert Emails** field, you can also specify extra email accounts that can receive alerts for this alarm.
- Go to the **Layout** tab and define columns, sort order, and the direction. The **Layout** tab only exists for **Expression** and **SQL Query** types. For **Script** type you have to provide a script to define the layout. After the layout is defined, you can control the data to be returned in the report. For example, if you want to include the current pending reasons for jobs that are pending longer than a day, then add the pending reason column to the layout.
- Go to the **Breached Items** tab and view all the breached items.

## Results

After the template is created, you can view it under **Alert Template Management**

## Creating an Alert

### About this task

Create an alert to notify you when certain conditions occur in your cluster.

### Procedure

- Go to **Alerts** in the Management section.
- Click **Add** on the upper right side.
- From the **Select Template** menu, either select **None** or any created templates.

**Note:** Created templates in the list are used to create alerts that are based on an existing template. If you do not see any templates in the list, you can go to **Alert Templates** in the Templates section and create a template. Alternatively, you can choose **None** and continue with alert instance.

- Click **Create**.

**Note:** If you want to change any alert values propagated from the template you can clear the **Template Propagation** check box.

## Results

Created alerts are displayed on the list of Alerts. You can stop or resume (if required), edit or delete the selected alert.

## Viewing Alert Log

### About this task

RTM keeps a log for all the alert activities.

### Procedure

1. Click the **Grid** tab.
2. Click **Alerts** on the left navigation pane.
3. Click **Alerts Log** tab. You can edit, view breached items, view alert history, enable or disable an alert.

---

## Syslogs alerts

Use **Syslogs** for viewing UNIX log file entries

Click the **Syslogs** tab to view UNIX log file entries in the **Syslogs** page.

The **Syslogs** page displays entries from the UNIX log files in the `/var/log` directory in each host in the clusters that RTM monitors. Each data record displayed here is an entry in one of the log files.

If you want the UNIX log files for a host to be displayed in Platform RTM, edit the `/etc/syslog.conf` file on that host and add the IP address of the Platform RTM host to the end of the file. For more information about configuring log files in the UNIX host, see the topic, “Enabling remote hosts syslog transfer,” in *Platform RTM Installation*.

You can create an alert rule (to notify you of future log entries) or a removal rule (to automatically remove log file entries) by clicking **Alerts** or **Removals** in the **Rules** window at the upper right of the **Syslogs** page.

The page contains the following fields:

- **Host.** The name of the host in which the log file entry is recorded.
- **Date.** The date of the log file entry, which is taken from its time stamp.
- **Time.** The time of the log file entry, which is taken from its time stamp.
- **Message.** The contents of the log file entry.
- **Facility.** The name of the system, log, or service that recorded the log file entry.
- **Level.** The error level of the log file entry.
- **Options.** Actions that you can perform on the log file entry. You can either create a removal rule or create an alert rule that is based on this entry.

---

## Chapter 5. External Load Indices manager (ELIM) Graphs

External Load Indices manager reports the values of dynamic external resources. A dynamic external resource is a user-defined resource with a numeric value that changes over time, such as the space available in a directory. Using RTM, you can create graphs for those ELIMs and analyze them.

If you have specific workload or resource requirements at your site, the LSF administrator can define external resources. The LSF administrator programs the elim to define external load indices, populate those indices with the values of dynamic external resources, and return the indices and their values to *stdout*.

The RTM administrator can create ELIM Templates based on ELIMs that were created for LSF, instantiate onto one or more clusters through Host Type, Host Model, or Host Group filter. The reason for this ability to filter hosts for ELIM Graph creation is that in many cases, an ELIM-based resource might be present only on a subset of hosts in a cluster.

---

### Creating an ELIM Template

#### Procedure

1. Go to **ELIM Templates** in the **Templates** section.
2. Click **Add** on the upper right side.
3. Enter all the required values to create an **ELIM Template** and click **Create**. After an **ELIM Template** is created, it can be instantiated onto one or more clusters through Host Type, Host Model, or Host Group filter.

---

### Creating an ELIM Instance

#### Procedure

1. Go to **ELIMs** in the **Management** section.
2. Click **Add** on the upper right side. Then, the **ELIM Instance Selection** tab shows up where you can enter the required values. If you created **ELIM templates**, then you can also import those templates and create an *ELIM instance*.
3. Click **Create**.

### Creating a Graph for an ELIM Instance

#### Procedure

1. Go to **ELIMs** in the **Management** section.
2. Select an *ELIM instance*. You can also select multiple instances at a time.
3. From the menu, select **Create Graphs Immediately** and proceed.
4. Click the *ELIM instance* and click the **Existing Graphs** tab.
5. Under **Existing Graph** tab, you can see a list of graphs that were created for that instance.
6. Click any graphs name under the **Graph title** column. You can see the details of each graph. You can even modify the values in the details section.

---

## GPFS ELIM

IBM General Parallel File System (GPFS™) is a high performance cluster file system. GPFS is a shared disk file system that supports the AIX®, Linux, and Windows operating systems. The main differentiator in GPFS is that it is not a clustered File System but a parallel File System. This means that GPFS can scale almost infinitely. Using RTM, you can monitor GPFS data.

In the RTM GUI, you can monitor GPFS on a per LSF host and a per LSF cluster basis either as a whole or per volume level.

Host level:

- Average MB In/Out per second
- Maximum MB In/Out per second
- Average file Reads/Writes per second
- Average file Opens/Closes/Directory Reads/Node Updates per second

Cluster level:

- Mb available capacity In/Out
- Resources can be reserved and used upon present maximum available bandwidth. For example, bsub to reserve 100 kbytes of inbound bandwidth at cluster level for 20 minutes: `bsub -q normal -R "rusage[gtotalin=100:duration=20]" ./myapplication myapplication_options`

## Configuring ELIM Script

### About this task

Configure the following ELIMs in LSF before proceeding:

- `elim.gpfshost` - Monitors GPFS performance counters at LSF host level
- `elim.gpfsglobal` - Monitors available GPFS bandwidth at LSF cluster level

Refer to the Administer LSF documentation for more details on configuring GPFS ELIMs. The ELIM Scripts are available from LSF 9.1.1 or later versions.

## Configuring LSF cluster

### Procedure

1. Add GPFS node as LSF server, computenode, or as master candidate.
2. Configure external load indices as LSF resources, for example:

|                                      |              |     |     |
|--------------------------------------|--------------|-----|-----|
| <code>gstatus</code>                 | String (30)  | ( ) | ( ) |
| <code>gbytesin</code>                | Numeric (30) | Y   | ( ) |
| <code>gbytesout</code>               | Numeric (30) | Y   | ( ) |
| <code>gopens</code>                  | Numeric (30) | Y   | ( ) |
| <code>gcloses</code>                 | Numeric (30) | Y   | ( ) |
| <code>greads</code>                  | Numeric (30) | Y   | ( ) |
| <code>gwrites</code>                 | Numeric (30) | Y   | ( ) |
| <code>grdir</code>                   | Numeric (30) | Y   | ( ) |
| <code>giupdate</code>                | Numeric (30) | Y   | ( ) |
| <code>gbytesin_gpfs_dev_name</code>  | Numeric (30) | Y   | ( ) |
| <code>gbytesout_gpfs_dev_name</code> | Numeric (30) | Y   | ( ) |
| <code>gtotalin</code>                | Numeric (30) | N   | ( ) |
| <code>gtotalout</code>               | Numeric (30) | N   | ( ) |

3. Map the external resources to hosts in the ResourceMap section of `lsf.cluster.cluster_name`. For example:

Begin ResourceMap



## RESOURCENAME LOCATION

- # GPFS Per Host Resources

```
gstatus ([hostgpfs01] [hostgpfs02] [hostgpfs03])
gbytesin ([hostgpfs01] [hostgpfs02] [hostgpfs03])
gbytesout ([hostgpfs01] [hostgpfs02] [hostgpfs03])
gopens ([hostgpfs01] [hostgpfs02] [hostgpfs03])
gcloses ([hostgpfs01] [hostgpfs02] [hostgpfs03])
greads ([hostgpfs01] [hostgpfs02] [hostgpfs03])
gwrites ([hostgpfs01] [hostgpfs02] [hostgpfs03])
grdir ([hostgpfs01] [hostgpfs02] [hostgpfs03])
giupdate ([hostgpfs01] [hostgpfs02] [hostgpfs03])
gbytesin_gpfs01 ([hostgpfs01] [hostgpfs02] [hostgpfs03])
gbytesout_gpfs01 ([hostgpfs01] [hostgpfs02] [hostgpfs03])
gbytesin_gpfs02 ([hostgpfs01] [hostgpfs02] [hostgpfs03])
gbytesout_gpfs02 ([hostgpfs01] [hostgpfs02] [hostgpfs03])
```

- # GPFS shared resources

```
gtotalin [all]
gtotalout [all]
```

End ResourceMap

4. Copy the elim executables to your cluster(\$LSF\_SERVERDIR). For example:

```
#cp elim.gpfs host elim.gpfs global $LSF_SERVERDIR
```

By default, the ELIM executable is stored in /opt/rtm/etc.

5. Reconfigure your cluster.

```
#lsfadmin reconfig
#badmin mbdrestart
```

For more information, see *LSF\_v9.1\_Doc\_Center/lsf\_admin/elim\_config\_enable\_lsf.html*.

## Creating graphs for GPFS ELIM

### Procedure

1. Log in to RTM, go to **Console->Management->ELIMs**
2. Click **Add** on the upper right side. Then, the **ELIM Instance Selection** tab shows up where you can enter the required values.
3. Select the template as **GPFS Bandwidth** or **GPFS File Operations**.
4. Select a Cluster Name and the Host Type and click **Create**.
5. Select a *GPFS ELIM instance*.
6. From the menu, select **Create Graphs Immediately** and proceed.



---

## Chapter 6. Managing the metadata plug-in

The metadata plug-in can add metadata for users, user group, host, host group, cluster, job group, queue, project, license-project, application, and queue group.

A few concepts to know:

### LSF object

Represents the various objects in LSF such as users, hosts, host groups, and queues.

### Meta configuration file

An XML file that outlines the attributes of each LSF object. It is like a blueprint for LSF objects. There can be only a single meta configuration.

### Metadata

A set of descriptive information for all attributes of each LSF object instance.

**Example:** A common use for the metadata plug-in is when you want to display metadata for a host, such as geographic location or rack location.

#### 1. Create metadata configuration file.

When you are defining XML configuration file, refer to the descriptions of metadata configuration file tags in “Define metadata XML configuration file.”

#### 2. Import the XML or CSV metadata file.

Access to the Metadata Settings pages is controlled in Realm Permissions. For more information, see “Enabling metadata settings in Realm Permissions” on page 65.

- Import configuration file through the command line.

### Examples:

XML configuration file:

```
php /opt/cacti/plugins/meta/metadata_cli.php -c -f <filename>
```

CSV file:

```
php /opt/cacti/plugins/meta/metadata_cli.php -l -t <object_type> -f <filename>
```

- To import configuration file through the RTM Console, see “Importing metadata XML or CSV configuration file via RTM Console” on page 65.

---

## Define metadata XML configuration file

By default, the sample metadata configuration file is at `/opt/cacti/plugins/meta/metadata.conf.xml`.

Refer to this list of metadata file tags when you are preparing the configuration file:

### <meta>

The enclosing tag for the entire metadata configuration file

**<object>**

Specifies the metadata object type to be configured; it can be any of the following types: user, user-group, host, host-group, cluster, job-group, queue, project, license-project, application, or queue-group

**<section>**

Specifies a new section; the on hover metadata window creates a tab for each section that can organize the data based on your preference

**<column>**

Configures a new metadata column

**<db\_column>**

The database column that stores this column's data (such as meta\_col1....meta\_col20)

**<display\_name>**

The label to be displayed for this column in the UI

**<description>**

The description to be displayed for this column in the UI; this value is displayed when you are adding a metadata object in the UI

**<data\_type>**

The type of data to be stored in this column; for most columns, the data type is text

Valid values are:

- Text: commonly used for most columns; the plug-in displays the value in the UI
- URL: when the data represents a hyperlink, the plug-in formats this column as a link
- Callback: for this value, the plug-in calls the URL provided and displays the result in the UI, without any parsing; dynamic data can then be presented as metadata like the current temperature of a host but it does have performance consequences. This tag must be used with caution.
- Cluster\_id: represents a cluster id; primarily used by queue groups
- Display\_name: represents a display name; primarily used by queue groups
- Queue\_list: represents a list of LSF queues; primarily used by queue groups

**<position>**

Indicates the order in which columns are displayed in the UI, ascending, starting with 1; this value must be unique for the type

**<summary>**

Indicates whether this column is displayed in the summary table for metadata object otherwise, the column is displayed only in the detail drill-down

**<search>**

Determines whether this column is searchable by the search condition that is specified in the metadata UI

**<popup>**

Determines whether this column is displayed in the metadata window for this object type

**Important:** The object\_id must exactly match the actual LSF object instance. For example, if 'wfeng' is specified as the object ID then it must match exactly this string in order for metadata to display the window.

---

## Define metadata using CSV file

You can choose to load all metadata information through a CSV file. The file must be in the following format:

- Header: [object\_id, cluster\_id, db\_column(s)] from grid\_metadata table.  
Essentially this information comes from the metadata configuration file.
- Data Rows: data values that follow the order of the header.

### Example:

```
object_id, cluster_id, meta_col1, meta_col2
wfeng, 1, William Feng, 9059484530
```

An example of a User object type, where meta\_col1 represents user's full name and meta\_col2 is the user's phone number is shown.

**Important:** The object\_id must exactly match the actual LSF object instance. For example, if 'wfeng' is specified as the object id then it must match exactly this string in order for metadata to display the window.

---

## Enabling metadata settings in Realm Permissions

### About this task

You must enable the metadata settings in **Realm Permissions** tab to access the metadata pages in the RTM Console.

### Procedure

1. Click the **Console** tab.
2. Under the Utilities section of the Console menu bar, click **User Management**.
3. On the user Management page, click the User Name that you want to grant metadata access.
4. Click **Realm Permissions** tab.
5. Click **RTM -> View Metadata** and **RTM -> Edit Metadata** check boxes to enable these settings.
6. Click **Save**.

The **Metadata Settings** link becomes available under the Configuration section of the Console menu bar.

---

## Importing metadata XML or CSV configuration file via RTM Console

### Procedure

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Metadata Settings**.
3. Click **Browse** to open the metadata configuration XML or CSV file you want to load.
4. Click **Import**.

Only XML or CSV files with well-defined metadata configuration are loaded.

- To edit an object, click the link from the list then click **Save**.

**Note:** You can organize queues into groups when you are editing the Queue Group entity. The Queue Group details are displayed on the Queue page when it is enabled in the Setting tab under the Queue section.

- To delete the metadata objects, select the metadata files on the list then choose **Delete** action before you click **Go**.

---

## Viewing metadata details

### Procedure

Hover the mouse over the name of an object.  
Popup window shows the metadata details.

---

## Chapter 7. Performance and maintenance

Performance and maintenance tasks include performance tuning for large clusters, database maintenance tasks, and setting up a remote database host.

---

### Performance tuning for large clusters

#### Store rrd files on a separate disk

##### About this task

To improve performance, store the rrd files on a separate disk from the database, and create a symbolic link to the new location.

##### Procedure

1. Create a directory on a separate disk.

For example:

```
mkdir /newDirectory
```

2. Copy files to your new directory.

For example:

```
cp -p /opt/cacti/rra/* /newDirectory
```

3. Make a backup of your existing files and create a symbolic link to your new directory.

For example:

```
mv /opt/cacti/rra /opt/cacti/rra.bak
ln -s /newDirectory /opt/cacti/rra
```

#### Increase the database memory

##### About this task

If you see errors that the database is running out of memory, increase the maximum memory that is allocated to the database.

MySQL 1114 errors usually indicate that the database needs more memory.

##### Procedure

1. Edit the `/etc/my.cnf` file and increase the `max_heap_table_size` value.
2. Restart the `mysqld` service:

```
service mysqld restart
```

#### Enable on-demand rrd file updating for systems with heavy disk I/O

##### About this task

If you have problems with high disk I/O wait times, enable on-demand rrd file updating. After it is enabled if you still have problems with high disk I/O wait times, you can also look into using Spine. It is an add-on Cacti feature.

## Procedure

1. Click the **Console** tab.
2. Under the **Configuration** section, click **Settings**.
3. Select the **Boost** tab.
4. In the **On Demand RRD Update Settings** section, select **Enable On Demand RRD Updating**.
5. You can optionally change how often your RRD files are updated by modifying the values for **How Often Should Boost Update All RRDs** and **Maximum Records** fields.
6. Click **Save**.

## Configure concurrent poller processes

### About this task

#### Procedure

1. Click the **Console** tab.
2. Under the **Configuration** section, click **Settings**.
3. Click the **Poller** tab.
4. In the **General** section, change the value for **Maximum Concurrent Poller Processes**.
5. Click **Save**.

## Enable database record partitioning

### About this task

Database record partitioning splits the larger LSF job data tables into multiple tables and speeds up processing during database maintenance operations.

Partitioning is required if you have many jobs per day or if you want to extend the time to keep job summary data.

It is recommended that you size the partitions to have a maximum of 2 million records. When you are sizing the partitions, consider the amount of memory the host has for the database. Each job record occupies 4 KB in the database.

You can also specify what elapsed time period to use for each partition. Increasing the time period means that the database contains more data for overall analysis, but also increases the system impact of removing job records.

#### Procedure

1. Click the **Console** tab.
2. Under the **Configuration** section, click **Grid Settings**.
3. Select the **Maint** tab.  
The **Grid Settings (Maint)** page is displayed.
4. In the **Large System Settings** section, select **Enable Record Partitioning** to place a check mark.
5. Specify elapsed time period in the **Partition Size** field.
6. Specify the maximum number of partitions to keep in the database in the **Maximum Partitions** field.



When the number of partitioned tables reaches this number, Platform RTM deletes the oldest partitioned table before it creates a new partitioned table. Increasing the number of partitions means that the database contains more data for overall analysis but also increases the size of the database.

7. Click **Save**.

## Configure data collection frequency

### About this task

For large clusters, it is recommended to change the data collection frequency. Data collection frequency is configured for each LSF cluster.

If you see continuous errors similar to the following in the `cacti.log` file, decrease the data collection frequency:

```
ERROR: Run-On/Abended Process Detected for ClusterName:'Large Cluster', ClusterID:'1',
Process:'GRIDJOBS', PID:'19749', Attempting to Kill PID
```

### Procedure

1. Click the **Console** tab.
2. In the **Grid Management** section, click **Clusters**.
3. Select the name of the LSF cluster you want to modify.
4. Click the **Poller** tab.
5. You can modify values for the following fields:
  - In the **Queue/Host/Load Collection Settings** section:  
**Collection Frequency** and **Max Allowed Runtime**
  - In the **Job Collection Settings** section: **Minor Collection Frequency**, **Major Collection Frequency**, and **Max Allowed Runtime**.
6. Click **Save**.

## Increase LSF API timeout values

### About this task

If you see errors in the `cacti.log` file that indicates the LSF APIs are timing out, increase the timeout value.

### Procedure

1. Click the **Console** tab.
2. In the **Grid Management** section, click **Clusters**.
3. Select the name of the LSF cluster you want to modify.
4. Click the **Advanced** tab.
5. In the **Cluster Connection Timeout Settings** section, modify values for the following fields:
  - **Base Timeout**
  - **Batch Timeout**
  - **Batch Job Info Timeout**
  - **Batch Job Info Retries**
6. Click **Save**.

## Enable Ajax filter

### About this task

Enable ajax filter in the **Grid Settings** page and enter at least three '%' characters in the **Grid** tab to get a list of hosts, host group, users, or user group.

### Procedure

1. Click the **Console** tab.
2. Under the **Configuration** section, click **Grid Settings**.
3. In the **General** tab, check **Enable Ajax filter** check mark.
4. Click **Save**.

---

## Maintaining Database

The following are the actions that you can take to maintain or optimize the database:

- “Set up record partitioning in the database”
- “Back up and restore the Platform RTM system settings” on page 71

## Set up record partitioning in the database

### About this task

In a large database, removing old job records and other database maintenance tasks can have a significant effect on database performance. You can improve database performance by enabling record partitioning. Partitioning divides tables within the database into multiple, smaller tables.

Removing job records from smaller tables has less of an impact on the system performance of active tables than removing job records from larger tables.

You can set up partitioning for the following database tables that Platform RTM monitors (all of the following tables are affected):

- GRID\_ARRAYS
- GRID\_JOB\_DAILY\_STATS
- GRID\_JOBS\_FINISHED
- GRID\_JOBS\_MEMPERF
- GRID\_JOBS\_RUSAGE

### Procedure

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings**.
3. Click the **Maint** tab.

The **Grid Settings (Maint)** page displays.
4. Under **System Maintenance Settings** section of the page, you can specify the time in **Database Maintenance Time** field to indicate when old database records must be removed from the database.

All partitioned data are always retained based on the retention periods you specified under the section of the page.
5. Scroll down to the **Large System Settings** section of the page, and ensure that the following options are set:

- **Enable Record Partitioning** - Check this box to enable record partitioning in your database. Partitioning is required if you have many jobs per day or if you want to extend the time to keep Job Summary data.
- **Partition Size** - Specify a time between different partitions.

Increasing the period means that the database contains more data for overall analysis, but also increases the system impact of removing job records.

6. Click **Save** to save your database record partitioning settings.

## Results

The partitioned tables have a numerical suffix (`_v000`, `_v001`, ... `_vnnn`). The newest table has the largest number.

## Back up and restore the Platform RTM system settings

Platform RTM saves system and graph settings in a MySQL database, which is the Cacti database. You can use the Platform RTM Console to back up operations of the Cacti database so that you can restore the Platform RTM system later. Having the backups means that if you changed the configuration in your Platform RTM host that you would like to revert (for example, if you incorrectly imported a Cacti template and your graphs no longer update, or you accidentally removed a cluster from the Platform RTM Console), you can use the Platform RTM Console to restore the Cacti database, which restores the Platform RTM system and graph settings and revert to your original, backed-up settings.

### Restriction:

The Platform RTM Console cannot restore all of the job performance data of the clusters that Platform RTM monitors. Any such data still in Platform RTM is likely out of date.

If you want to fully back up the entire database, you must use the **mysqldump** command. In addition, you cannot use the MyISAM table structure if you use the **mysqldump** command regularly. For more information, see MySQL documentation (for example, refer to the MySQL website at <http://www.mysql.com>).

The following is a list of actions you can take to back up or restore the Platform RTM system and graph settings through the Platform RTM Console:

- “Back up the Platform RTM system using the Platform RTM Console” on page 72
- “Enable automatic data archiving” on page 73
- “Restore the Platform RTM system using the Platform RTM Console” on page 73

The following files are backed up:

- `lsfpollerd.conf`: Database file that contains the credentials
- `lsf.conf`: The `lsf.conf` file that is associated with each cluster.
- `ego.conf` (for LSF 7.x and LSF 8.0 clusters only): The `ego.conf` file that is associated with each cluster.
- `server.key` and `server.crt`: These files are used as certification to `httpd` for `advocate`.

All tables in the Cacti database except the following list are backed up:

- `GRID_ARRAYS`

- GRID\_ARRAYS\_FINISHED
- GRID\_JOB\_DAILY\_STATS
- GRID\_JOB\_INTERVAL\_STATS
- GRID\_JOBS
- GRID\_JOBS\_FINISHED
- GRID\_JOBS\_JOBHOSTS
- GRID\_JOBS\_JOBHOSTS\_FINISHED
- GRID\_JOBS\_MEMPERF
- GRID\_JOBS\_PENDREASONS
- GRID\_JOBS\_PENDREASONS\_FINISHED
- GRID\_JOBS\_REQHOSTS
- GRID\_JOBS\_REQHOSTS\_FINISHED
- GRID\_JOBS\_RUSAGE
- LIC\_DAILY\_STATS
- LIC\_FLEXLM\_LOG
- LIC\_INTERVAL\_STATS
- LIC\_LUM\_EVENTS
- POLLER\_OUTPUT
- POLLER\_OUTPUT\_BOOST
- SYSLOG
- SYSLOG\_REMOVED

## Back up the Platform RTM system using the Platform RTM Console

### About this task

Before you upgrade to a newer version of Platform RTM, you must first back up the existing Cacti database.

Backing up the Cacti database is also recommended during scheduled server maintenance and you can revert to previous Platform RTM system and graph settings.

### Procedure

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings**.
3. Click the **Maint** tab.
4. Scroll down to the **Database Backups** section of the page, and ensure that the following options are set (if some of them are not needed, you can leave them blank):
  - Backup Cacti Database - Check this box to ensure that the Cacti database is backed up when the maintenance script runs.
  - Backup Schedule - Specify backup schedule as Daily or Weekly
  - Weekly Backup Day - If you chose Weekly backup, specify the day that you want the backup to take place
  - Backup Generations - Specify the number of backup file copies you want to maintain

- Database Backup Location - If a backup directory does not yet exist, provide a location. (If the directory exists, the message “[OK: DIR FOUND]” displays under the directory field.)
  - Post Backup Command - Provide the post backup script to be run after the Cacti database and the scheduled backup are completed. Ensure that the file is both executable and readable by the RTM service account.
5. Under the Utilities section of the Console menu bar, click **Grid Utilities**.
  6. In the Database Administration section of the **Grid Utilities** page, click **Force Cacti Backup**.  
The **Backup Files** table at the bottom of the page updates with the newly created backup file, modification date, and file size.
  7. Click the backup file name to download it to a specified location.
  8. Verify that the .tgz file contains the following files after the download is complete:
    - cacti\_db\_backup.sql
    - cacti\_db\_struct\_backup.sql
    - rtm/etc/rtm.lic
    - rtm/etc/lspollerd.conf
    - rtm/etc/server.crt
    - rtm/etc/server.key
    - rtm/etc/cluster\_id/lsf.conf
    - rtm/etc/cluster\_id/ego.conf (for LSF 7.x and LSF 8.0 clusters only)

After all files are successfully verified and backed up, you can upgrade to a new Platform RTM version or perform server maintenance without fear of losing or corrupting your existing database.

## Enable automatic data archiving

### About this task

Enable data archiving to save legacy job and job-related data to an archive database during scheduled server maintenance, and to archive job detail records to an archive directory or file server.

### Procedure

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings**.
3. Click the **Archiving** tab.
4. To enable data archiving of legacy job and job-related data, select the **Enable Data Archiving** box and specify the data archiving settings.
5. To enable data archiving of job detail records before data purging, select the **Create RRD's During Job Detail Purge** box and specify the path to the archive directory or file server.
6. Click **Save** to save your data archiving settings.

## Restore the Platform RTM system using the Platform RTM Console

### About this task

After you complete any scheduled server maintenance, or finish upgrading your Platform RTM version, you must restore the Cacti database that you previously backed up.

You can also restore the Cacti database to revert to a previous Platform RTM system and graph setting.

### Procedure

1. Click the **Console** tab.
2. Under the Configuration section of the Console menu bar, click **Grid Settings**.
3. Click the **Maint** tab.
4. Scroll down to the Database Backups section of the page, and find the **Database Restore** option.
5. Browse to the location of your previously backed-up database, and then click **Save** to upload and restore the .tgz or .zip backup file.

If the file is successfully restored, the message "Save successful" displays.

### Manually back up or restore the database Platform RTM

For later versions, all components are backed-up and packed into a .tgz file by the Platform RTM Console. However, previous versions of Platform RTM require that you manually back up your database. The following steps describe the manual backup and upgrade process:

1. "Back up the website"
2. "Back up the MySQL database"
3. "Back up the Platform RTM poller" on page 75
4. "Upgrade the Cacti website and grid plug-in" on page 75
5. "Upgrade the MySQL database" on page 76
6. "Upgrade the Platform RTM poller" on page 76

#### Back up the website:

This backup is recommended whenever you upgrade Cacti. You must back up all directories except the log and rra directories. If you are not concerned about the size of the backup, you can continue to back up those directories though. The following is a command example:

```
cd /var/www/html
tar --exclude=*.rrd --exclude=*.log -cf cacti_backup.tar cacti
```

If your Platform RTM upgrade does not involve the Cacti product, you can greatly speed and simplify your backup by using the following command:

```
cd /var/www/html/plugins
tar -exclude=*.rrd -cf grid_backup.tar grid
```

#### Back up the MySQL database:

Platform RTM provides a utility for backing up your Platform RTM database. When used, it backs up your database schema and important Cacti and Platform RTM tables to disk. It does not back up some rather large tables in Platform RTM because of their size and purpose. The Cacti/Platform RTM tables that are not backed up include:

- GRID\_ARRAYS
- GRID\_ARRAYS\_FINISHED
- GRID\_JOB\_DAILY\_STATS
- GRID\_JOB\_INTERVAL\_STATS
- GRID\_JOBS

- GRID\_JOBS\_FINISHED
- GRID\_JOBS\_JOBHOSTS
- GRID\_JOBS\_JOBHOSTS\_FINISHED
- GRID\_JOBS\_MEMPERF
- GRID\_JOBS\_PENDREASONS
- GRID\_JOBS\_PENDREASONS\_FINISHED
- GRID\_JOBS\_REQHOSTS
- GRID\_JOBS\_REQHOSTS\_FINISHED
- GRID\_JOBS\_RUSAGE
- LIC\_DAILY\_STATS
- LIC\_FLEXLM\_LOG
- LIC\_INTERVAL\_STATS
- LIC\_LUM\_EVENTS
- POLLER\_OUTPUT
- POLLER\_OUTPUT\_BOOST
- SYSLOG
- SYSLOG\_REMOVED

**Note:**

SSH private key pair files for grid cluster control authentication are not backed up; the backup process backs up only the file path to the key file.

Even though the contents of the tables are not backed up, their structure is. To run the database backup, you must first specify a backup location under **Grid Settings** > **Maint** (in the Configuration section of the Console menu bar). After you verify that the backup is complete, open a shell and do the following steps:

```
cd /var/www/html/cacti/plugins/grid
php -q database_backup.php
```

Running that utility creates three backup files: the MySQL database itself, the Cacti/Platform RTM Schema, and Critical Platform RTM Tables. These backup files can be used in the case of a disaster to restore full functionality to the system.

**Back up the Platform RTM poller:**

To support the ever growing list of features and to provide for bug fixes, the Platform RTM pollers must be upgraded from time to time. Therefore, before you upgrade Platform RTM, first make a backup copy of your Platform RTM pollers.

To back up your pollers, run the following commands on each poller, or back up one poller and distribute the changes to all other pollers.

```
cd /opt/rtmtar -cf grid_poller_backup.tar bin lic 1sf61 1sf62 1sf701 1sf702 1sf703 1sf704 1sf705 1sf706 1sf8
```

**Upgrade the Cacti website and grid plug-in:**

To upgrade the Cacti website, you must pay particular attention to a few points:

1. You must make a backup copy of your <path\_to\_cacti>/include/global.php file. This file is important as it contains your database connect string and your list of installed plug-ins.
2. You must know whether you modified the default Platform RTM poller scripts. If you modified them, you must make a separate backup of your

<path\_to\_cacti>/scripts directory. If you added new scripts that are not a part of the standard Platform RTM deployment, it is not necessary to back up this directory.

3. If you modified default Platform RTM Data Queries, you must back up your <path\_to\_cacti>/resource directories. Again, if you did not modify the default installation of Platform RTM Data Queries, it is not necessary to proceed with this step.

After completing these steps, note the ownership of all files in Platform RTM. The Platform RTM installer recommends that the cacti user account own the Apache site, and that the Apache binaries be run by the same. If so, then all files in the /var/www/html/cacti directory and subdirectories are “cacti:cacti” regardless of file ownership, not the ownership and permissions.

When you receive your Platform RTM upgrade, there are two directories in the upgrade path. They include both the “cacti” and “poller” directories. To upgrade the Cacti Website and Grid Plug-in, run the following command:

```
cp -rp <upgrade_path>/cacti /var/www/html
chown -R cacti:cacti /var/www/html
```

Verify that your config.php remains unchanged. Then, refresh your browser that is already pointed to the Platform RTM website.

### **Upgrade the MySQL database:**

After you upgrade your Cacti Website and Grid Plug-in, you must run a Database Upgrade script. As Platform RTM matures, several changes are made to the Platform RTM tables to accommodate new features.

#### **Note:**

As of Platform RTM 1.01, larger tables such as the grid\_jobs and grid\_jobs\_usage do not typically undergo structural changes. If one of these larger tables were to require modification, you are informed up-front to advise you of the downtime that might be required to upgrade any of those tables.

To upgrade your database schema, run the following commands from a shell:

```
cd /var/www/html/cacti/plugins/grid
php -q database_upgrade.php
```

### **Upgrade the Platform RTM poller:**

1. Verify the permissions on all file objects within the Platform RTM poller directory.

They are typically owned by “cacti:cacti”. If you chose another method, note those file permissions.

2. Run the following commands to upgrade the pollers:

```
cp -rp <upgrade_path>/poller /usr/local
chown -R cacti:cacti /usr/local/grid
```



---

## Setting up a remote database host

### Before you begin

You must ensure that the MySQL database server port (3306) is open on the firewall. In addition, the intended database host must meet the following requirements:

- The host is running a supported operating system.  
Platform RTM 9.1 supports Red Hat Enterprise Linux (RHEL) 5.2 and above or CentOS 5.2 and above.
- The host has at least 4 GB of RAM and 40 GB of hard disk space.
- The host has MySQL Server installed.
- The database host name and the Platform RTM host name can resolve to each other.

### About this task

By default, the Platform RTM host runs as the database host. You can choose to set up the database on another host to disperse Platform RTM activity to other machines, or to use a more optimal database host.

#### Important:

Go through these steps during the initial setup of your Platform RTM installation. If you choose to do these steps after Platform RTM is in operation for a while, the database can take a significant time to transfer to another host, during which the database is down.

### Procedure

1. Manually back up the existing Platform RTM database by using the Platform RTM Console.

Follow the steps that are described in “Back up and restore the Platform RTM system settings” on page 71 and “Back up the Platform RTM system using the Platform RTM Console” on page 72.

2. Prepare the Platform RTM host for the database transition.

- a. Log in to the Platform RTM host.

- b. Stop all essential services by running the following commands:

```
service lsfpollerd stop
service crond stop
chkconfig mysqld off
```

- c. Create a SQL database memory dump by running the following commands:

```
mysqldump cacti > cactidump.sql
mysqldump mysql time_zone_name > timezone.sql
```

where *rtm\_host\_ip* is the IP address of the Platform RTM host and *time\_zone\_name* is the name of your time zone.

Record the current directory (run `pwd`), as you need this file path to copy the SQL database dump files to the new database host.

- d. Change the Platform RTM configuration files to refer to the new database host instead of the local host.

In the following files, go to the line where the database host name is being declared and change “localhost” to the host name of the new database host or change “127.0.0.1” to the IP address of the new database host:

- /opt/cacti/include/config.php
  - /opt/rtm/lsfversion\_number/bin/grid.conf
  - /opt/rtm/etc/lsfpollerd.conf
  - /etc/rsyslog.conf
  - /opt/cacti/plugins/syslog/config.php
  - /opt/rtm/etc/lic.conf
3. Set up the new database host for Platform RTM.
    - a. Log in to the new database host.
    - b. Set up the MySQL database services by running the following commands:
 

```
yum install mysql-server mysql
chkconfig mysqld on
/etc/init.d/mysqld start
```
    - c. Create a new Platform RTM database by using the following command:
 

```
mysqladmin -u root create cacti
```
    - d. Assign the appropriate database privileges by running the following commands:
 

```
mysql -u root mysql -e "GRANT ALL ON cacti.* TO cacti@rtm_host_ip IDENTIFIED BY 'admin';"
mysql -u root mysql -e "GRANT SELECT ON mysql.time_zone_name TO cacti@rtm_host_ip IDENTIFIED BY 'admin';"
```

where *rtm\_host\_ip* is the IP address of the Platform RTM host and *time\_zone\_name* is the name of your time zone.
    - e. Back up the original my.cnf file and copy the my.cnf.\*.innodb files from /opt/rtm/share in the Platform RTM host to /etc in the new database host by running the following commands:
 

```
cp my.cnf my.cnf.original
scp root@rtm_host_ip:/opt/rtm/share/my.cnf.*.innodb /etc/my.cnf
```

where *rtm\_host\_ip* is the IP address of the Platform RTM host.
    - f. Copy the SQL database dump files from the Platform RTM host to the new database host by running the following commands:
 

```
scp root@rtm_host_ip:path_to_sql_dump/cactidump.sql /tmp/
scp root@rtm_host_ip:path_to_sql_dump/timezone.sql /tmp/
```

where *rtm\_host\_ip* is the IP address of the Platform RTM host and *path\_to\_sql\_dump* is the file path to the database dump files.
    - g. Import the SQL database dump files into the new database by running the following commands:
 

```
mysql cacti < /tmp/cactidump.sql
mysql mysql < /tmp/timezone.sql
```
    - h. Restart the MySQL database by running the following command:
 

```
service mysqld restart
```
  4. Restart the services and verify that the remote database is set up properly.
    - a. Log in to the Platform RTM host.
    - b. Restart the essential services by running the following commands:
 

```
service crond start
service rsyslog restart
service lsfpollerd start
service licpollerd restart
```
    - c. Verify that the database is set up properly by running the following command:
 

```
service lsfpollerd status
service licpollerd status
```

If **lsfpollerd** and **licpollerd** are running, the database is set up properly.

---

## Notices

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

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

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

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan, Ltd.  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web

sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Intellectual Property Law  
Mail Station P300  
2455 South Road,  
Poughkeepsie, NY 12601-5400  
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application

programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. \_enter the year or years\_.

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

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)<sup>®</sup> are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

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



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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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







Printed in USA

SC14-7577-00

