

---

# Platform RTM Administrator Guide

Platform RTM  
Version 8.0.2  
Release date: October 2011



Copyright

© 1994-2011 Platform Computing Corporation.

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

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

We’d like to hear from you

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

Your comments should pertain only to Platform documentation. For product support, contact [support@platform.com](mailto:support@platform.com).

Document redistribution and translation

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

Internal redistribution

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

Trademarks

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

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

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

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

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

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

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

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

Third-party license agreements

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

---

# Contents

1	About Platform RTM .....	5
	Introduction to Platform RTM .....	6
2	Monitor clusters using Platform RTM .....	9
	Configure cluster management .....	10
	Configure cluster interaction .....	16
	Control an LSF cluster .....	21
	Enforce cluster action control comments .....	23
	Configure idle job detection for cluster .....	24
	Configure idle job detection (Grid settings) .....	25
	Configure memory violation detection .....	26
	Add or edit LSF clusters for Platform RTM to monitor .....	27
	View LSF cluster and job information .....	30
	View UNIX log file entries .....	48
	View time series graphs for clusters and hosts .....	49
	View non-time series charts for clusters and hosts .....	51
3	Working with application license servers .....	53
	Relationship between Platform RTM and license servers .....	54
	License Management section .....	55
	Configure license pollers .....	57
	Add or edit license servers for RTM to monitor .....	58
	Monitor license servers .....	59
	View time series graphs for license servers .....	62
	View non-time series charts for license servers .....	63
4	Working with thresholds and alerts .....	65
	Configure thresholds and alerts .....	66
	View threshold and alert information .....	68
	Acknowledge triggered alerts for thresholds .....	70
5	Metadata Plugin .....	71
	Manage metadata plugin .....	72
6	Generic alarms (advanced feature) .....	75
	Configure alarms .....	76
7	Performance and maintenance .....	81
	Performance tuning for large clusters .....	82
	Database maintenance .....	85

Set up a remote database host .....	92
Issues to consider .....	94

C H A P T E R

---

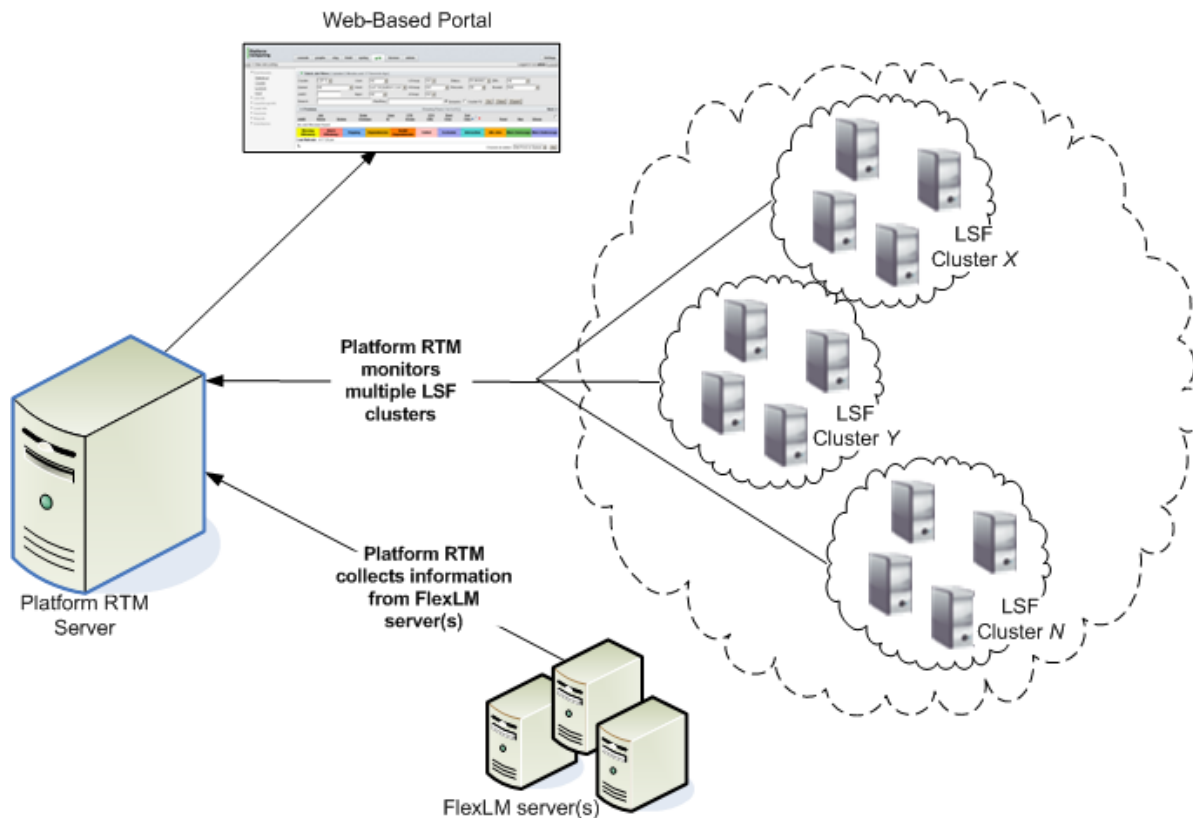
# 1

## About Platform RTM

# Introduction to Platform RTM

Platform RTM is an operational dashboard for 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 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.



Platform RTM is used to monitor and graph LSF resources (including networks, disks, applications, etc.) in a cluster. In graph or table formats, Platform RTM displays resource-related information such as the number of jobs 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 used by Platform RTM for license control. FLEXlm allows licenses to reside on the network instead of a specific host; this allows any registered host in the cluster to access and use a specific licensed application, thereby increasing usage efficiency.

## Platform RTM and Cacti

Cacti is a complete RRDTool-based graphing solution developed by The Cacti Group. 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 plugin so that

when used together, Platform RTM can offer LSF-specific monitoring and reporting capabilities 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, Platform provides specific data pollers (Platform data pollers) that work with Platform RTM. You are free to use your own data pollers to work with RTM, but Platform will not provide support these custom data pollers.

Platform RTM is licensed under the GNU General Public License (GPL) v2, and you can obtain the source code from the Platform FTP site. The Platform data pollers are licensed under the standard Platform End User License Agreement (EULA).

## About this guide

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 guide 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 <http://cacti.net/documentation.php>.

Platform also provides Platform RTM download, installation, and release information on [my.platform.com](http://my.platform.com).





## Monitor clusters using Platform RTM

- [Configure cluster management](#) on page 10
- [Configure cluster interaction](#) on page 16
- [Control an LSF cluster](#) on page 21
- [Enforce cluster action control comments](#) on page 23
- [Configure idle job detection for cluster](#) on page 24
- [Configure idle job detection \(Grid settings\)](#) on page 25
- [Configure memory violation detection](#) on page 26
- [Add or edit LSF clusters for Platform RTM to monitor](#) on page 27
- [View LSF cluster and job information](#) on page 30
- [View UNIX log file entries](#) on page 48
- [View non-time series charts for clusters and hosts](#) on page 51
- [View time series graphs for clusters and hosts](#) on page 49

# Configure cluster management

Click the Console tab and use the Console menu bar on the left side to configure cluster management.

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

The following is a list of actions that you can take to configure cluster management:

- Configure settings to monitor your cluster using Platform RTM.  
Navigate to the Management section of the Console menu bar. For more details on the pages in this section, see [Management section](#) on page 10.
- Add LSF clusters and perform certain database administration functions.  
Navigate to the Grid Management section of the Console menu bar. For more details on the pages in this section, see [Grid Management section](#) on page 13.
- Monitor license servers and pollers.  
Navigate to the License Management section of the Console menu bar.
- Grid Management: Pollers, Clusters, and Utilities pages  
These pages allow you to add LSF clusters and perform database administration functions.
- Use Cacti utilities and features.  
Navigate to the Collection Methods, Templates, or Import/Export sections of the Console menu bar to use Cacti utilities and features. For more details on these sections, refer to the Cacti documentation (<http://cacti.net/documentation.php>).

## Management section

The Management section is located in the Console menu bar.

### Thresholds page

Navigate to the Thresholds page by clicking Thresholds under the Management section of the Console menu bar. This page shows the configured thresholds in your cluster. 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, this is the amount of time from when the alert was first triggered.

- **Repeat.** The amount of time that the threshold waits before repeating the alert 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 has triggered an alert
- **Enabled.** Indicates whether this threshold is currently active.
- **Ack.** Indicates whether the threshold alerts have been acknowledged: "on" indicates that the threshold has been acknowledged; "off" indicates that the threshold either has not been acknowledged, or had its acknowledgement reset.

## Threshold Item

Navigate to the Threshold Item page for a threshold by clicking the name of the threshold from the Thresholds page. This page allows you to configure threshold settings and event triggering.

Event triggering behavior is based on re-alert cycle settings. When the threshold first triggers an alert, the event trigger is invoked based on a high or low threshold breach. If the alert stays triggered, the event trigger is invoked again unless the re-alert 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 this page:

- **Template propagation enabled:** Enable the propagation of changes to the threshold template
- **Threshold name:** The name of the threshold as it appears 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 has returned to normal
- **Reset acknowledgement:** Reset acknowledgements when the threshold has returned 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 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 additional information. Placeholders for the email message body are enclosed by angle brackets (<>), for example, <cluster\_name>.

- Syslog settings
- Data type: Special formatting for the given data
- Re-alert 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 based on your system and you can insert placeholders in alert email templates to present additional information for administrators to 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:

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.
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 linked in this alert.
Host description	host_description	<host_description>	The host description of the device linked in this alert.
Threshold description	Not available	<DESCRIPTION>	The threshold description.
Threshold host name	Not available	<HOSTNAME>	The host name of the threshold.

Placeholder name	Tag for threshold name	Tag for alert email template	Description
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 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 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 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 located in the Console menu bar.

### Pollers page

Navigate to the Pollers page by clicking Pollers under the Grid Management section of the Console menu bar. This page shows information about Platform RTM pollers. 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 new 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.
- **LSF Version.** The LSF version running on the associated cluster.
- **License Threads.** The number of license threads that the poller uses for data collection. Data collection is faster if you specify more license threads.
- **Physical Location.** The physical directory location of the local Platform RTM poller (for example, /opt/rtm/lsf version/bin). If the directory is found and verified, the message [OK: DIR FOUND] appears 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 for the pollers on which to perform an action (for example, Delete), then select an action and click go.

## Clusters page

Navigate to the Cluster page by clicking Clusters under the Grid Management section of the Console menu bar. This page shows information about LSF clusters (including configured time out thresholds, and job efficiency information) and the pollers that collect data from them.

For more information on configuring a cluster using Platform RTM, see [Add or edit LSF clusters for Platform RTM to monitor](#) on page 27.

- **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 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 performing Database maintenance.
  - **Admin Down.** This status appears if you unchecked the Should Daemons be Enabled option in the General Settings tab for the Poller.
  - **Up.** Indicates that there are queues, hosts, and loads collected within 3\*non-job-interval and job information collected within 3\*job-interval. These 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 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 the following: 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 could 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 CPUs 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 time out.
- **Job Minor Freq.** The configured job minor frequency.
- **Job Major Freq.** The configured job major frequency.
- **Job Timeout.** The configured job time out.
- **LIM Timeout.** The configured lim time out.
- **Choose an action.** Select one or more checkboxes for the clusters on which to perform an action (for example, Enable or Disable), then select an action and click go.

## Utilities page

Navigate to the Utilities page by clicking Grid Utilities under the Utilities section of the Console menu bar. This page shows 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.

- **View Grid Process Status.** Click to open the Grid Process Status page and show status information associated with cluster polling processes (for example, statistics for cluster poller, runtime, database maintenance, license collection, etc.).
- **Force Cacti Backup.** Click to perform a backup of key Cacti and Platform RTM database tables. See [Database maintenance](#) on page 85 for more information on database backup and restore.
- **Purge Completion Factor Data.** Click to purge the completion factor records from the tables, and allow for re-creation of completion factor information based upon a change in settings.
- **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. This section will not appear until Platform RTM creates at least one backup file.

## Configure cluster interaction

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

The Console menu bar has eight sections: Create (graphs), Management, Grid Management, Collection Methods, Templates, Import/Export, Configuration, and Utilities. Most are default Cacti utilities and features, and are documented in the Cacti documentation (<http://cacti.net/documentation.php>).

The following is a list of actions you can take to configure cluster interaction with Platform RTM:

- Configure how the clusters interact with Platform RTM.

Click Grid Settings under the Configuration section of the Console menu bar. For more details on the subtabs in this page, see [Grid Settings page](#) on page 16.

- Configure how Platform RTM interacts with the clusters.

Click Cacti Settings under the Configuration section of the Console menu bar. For more details on the subtabs in this page, see [Cacti Settings page](#) on page 19.

## Grid Settings page

Navigate to the Grid Settings page by clicking Grid Settings under the Configuration section of the Console menu bar.

Contact Platform Computing for assistance in determining the optimal settings for your clusters. 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 this page displays, take note of the following specific fields:

- Strip Domains

Makes the name output shorter for hosts belonging to the specified domain names. It allows you to conserve display space if you are displaying hosts in common domains within your organization.

- Summary Host name Substitute

Allows you to 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. You should 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. You should restrict this setting to reduce load because the job zoom function is system-intensive. You should 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 belonging to the group. The Job Specificat ion option only displays the jobs submitted with bsub - G option.

To set the method that matches your configuration, do the following:

1. Click the Console tab.



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

To see the display, navigate to the Grid tab then click Groups under User/Group Info section of the Grid menu bar. Click on 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 exported will reduce system performance when performing 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.

- Multiple Column Sort

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

---

### Note:

- Upon initial access to a page, a pre selected column for sorting is displayed. The display varies depending on the page you are viewing.
 

To add a sort column, click on the column name. The first column 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 on 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 on the **x** beside the column name. Once a sorting column is deleted, the rest of the columns automatically shift upward.
- License and Grid Plugin pages now support multiple column sort capability.

- All Platform RTM pages support the multiple column feature except for these 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, navigate to **Console** tab then click **User Management** under Utilities section of the Console menu bar. Select the **User** then click on **Realm Permissions** tab. Click the **RTM -> Multiple Column Sort** checkbox.

- 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 this page displays, take note of the following specific fields:

- Detail Job Data Retention Period

Records of job details are kept for this period of time after the job is ended. 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 of time after the job ended. The size of each job record depends on job volume and your cluster settings.

- Daily Statistics Retention Period

Record of daily summary statistics are kept for this period of time after the job ended. As these records are added every day, you can keep records for a longer period of 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 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 deep-dive analysis that will not affect the system, because you can perform this analysis on the archive database instead of on a database that is currently in use.

- Paths: Configures cluster directories and file paths. If the directories and file paths are found and verified, the message [OK: FILE FOUND] appears 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 should be treated when searching 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 menu bar. If you choose In Service, jobs with RES-Down status will not appear 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 this 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?**  
Enabling this allows you 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 Job Memory Resources be Tracked?**

If enabled, Platform RTM maintains an internal table of memory performance statistics.

---

**Note:**

Only enable this setting if you know how to access and use the internal table.

---

- **Aggregation Level**  
When you build a job group hierarchy, only a number of fields are significant for analysis. RTM only aggregates job groups by the specified level, if you set up a non-zero value.
- **Status/Events:** Configure default behavior, thresholds, and visual cues for job flapping, cluster and job efficiency, PID levels, and job dependencies.
- **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.
- **Memory Exceptions:** Configures detection and email notification settings for memory rusage violations.

## Cacti Settings page

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

- **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] appears 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 this 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 ] appears below this field.
- **Misc: Cacti Log Viewer for User**
- **Alerting/Thold:** Configure alert and cluster threshold settings.
- **Syslog:** Configure Syslog event settings.
- **RTM Plugins:** Configure Platform RTM plugin settings.

Among the fields that this 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 a lot of 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 (unchecked), 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.

# Control an LSF cluster

Platform RTM allows you to control LSF clusters, hosts, queues, and jobs, as long as you enabled Platform RTM to control the LSF clusters. Platform RTM controls the LSF clusters by invoking LSF commands in the LSF master host.

Platform RTM allows you to control the following LSF components:

- 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 reconfi g` (dynamically reconfigures LSF)
- `lsadmin reconfi g` (restarts LIM on all hosts in the cluster)

- Hosts

You can open or close hosts in LSF clusters.

- Queues

You can perform 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:

- `bt op` (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)
- `bswi tch` (switches unfinished jobs from one queue to another)
- `bresume` (resumes suspended jobs)
- `brun` (forces jobs to run immediately)
- `bst op` (suspends unfinished jobs)
- `bki ll` (sends signals to kill unfinished jobs)
- Force kill (forces a job kill using the `bki ll - r` command)
- Signal kill (sends a specific signal to kill a job using the `bki ll - s` command)

## Run grid control commands on an LSF cluster

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

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

1. From the Platform RTM console, click the Console tab.
2. Under the Grid Management section of the Console menu bar, click Clusters.

## Monitor clusters using Platform RTM

3. Click the Cluster Name link of the cluster you want to configure.
4. Configure the cluster settings. If the grid control requires additional information, specify the details in the displayed fields.
5. Click Save.

To see the status of a cluster, refer to [Dashboards section](#) on page 30.

# Enforce cluster action control comments

Platform RTM allows you to 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, etc.) for host, queue, or cluster can be submitted.

## Set mandatory cluster action control comments

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 checkbox.
5. Click Save.

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

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 drop-down list.
5. Click Go.

A confirmation dialog box appears.

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

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

---

**Note:**

After a fresh installation of RTM 8.0, your any previous settings for idle jobs are no longer in effect. You must configure them again.

---

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)

1. Click the Console tab.
2. Under the Configuration section of the Console menu bar, click Grid Settings and navigate 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

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

1. Click the Console tab.
2. Under the Configuration section of the Console menu bar, click Grid Settings and navigate 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 should be 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 prior to flagging the job.
  - 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 prior to flagging the job.
  - 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 need to take, in chronological order, to add or edit LSF clusters for Platform RTM to monitor:

1. [Add or edit clusters](#) on page 27
2. [Add the RTM host to the LSF cluster as an LSF client](#) on page 28
3. [Add all hosts in the LSF cluster to RTM as devices](#) on page 29

## Add or edit clusters

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

- [Add or edit clusters using the Platform RTM Console](#) on page 27  
You can use the Platform RTM Console to add or edit an LSF cluster for monitoring.
- [Add clusters to RTM using a script](#) on page 28  
You can use a script to add an LSF cluster for monitoring by Platform RTM.

## Add or edit clusters using the Platform RTM Console

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

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 a new 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 describing 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 drop-down list of Grid Poller.

5. Click Create (or Save) to save the settings for your LSF cluster.
  - 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 28.

## Add clusters to RTM using a script

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

1. From the command line, navigate 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 representing 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.

- Navigate to the Console tab then click Pollers under Grid Management section. The RTM Pollers page opens displaying the Poller Names and Poller ID. You should 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` located in `/share/lsf/conf`:

```
php -q grid_add_cluster.php --type=0 --pollerid=8 --cluster_name=maincluster --cluster_env=/
share/lsf/conf
```

After adding clusters to RTM using a script, you need to do the following:

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](#) on page 28.
2. Verify that the new cluster is added to RTM using 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

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

1. Log into 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 `mbatd` 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 into the RTM host.
  - b) From the RTM host, use `telnet` to log into 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.

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

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](#) on page 29.

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

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

1. In the command line, navigate 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 using the Grid Host template.

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

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 on configuring the interaction between Platform RTM and LSF, see [Configure cluster interaction](#) on page 16.

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.

Navigate to the Dashboards section of the Grid menu bar. For more details on the pages in this section, see [Dashboards section](#) on page 30.

- Control the cluster.

Navigate to the Management section of the Grid menu bar.

These pages allow you 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 on LSF jobs.

Navigate to the Job Info section of the Grid menu bar.

These 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 details on the pages in this section, see [Job Info section](#) on page 32.

- Review information on LSF users and user groups.

Navigate to the User/Group Info section of the Grid menu bar. For more details on the pages in this section, see [User/Group Info section](#) on page 41.

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

Navigate to the Load Info section of the Grid menu bar. For more details on the pages in this section, see [Load Info section](#) on page 42.

- Review information on LSF hosts and host groups.

Navigate to the Host Info section of the Grid menu bar. For more details on the pages in this section, see [Host Info section](#) on page 43.

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

Navigate to the Reports section of the Grid menu bar. You can also filter statistics and batch system parameters for specific information in these pages. For more details on the pages in this section, see [Reports section](#) on page 45.

- 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 details, see [Settings tab \(from the Grid page\)](#) on page 46.

## Dashboards section

The Dashboard section is located 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 will be changed to reflect the current filtering.

## 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 started during the last hour.
- **Hourly Done:** The total number of jobs completed during the last hour.
- **Hourly Exit:** The total number of jobs aborted 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.




This page shows the following status information represented by icons:

**Table 1: LIM Status**

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>

Icon	Status
	<ul style="list-style-type: none"> <li>• unavail</li> <li>• unlicensed</li> <li>• sbat chd is down</li> <li>• RES is down</li> </ul>

**Table 2: Batch Status**

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>
	<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.

## 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 appear as either small or large in size. Click the Settings tab and modify the settings found under the Visual sub-tab to control this behavior.

## Statistical Dashboard

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

## License Dashboard

The license link on the dashboard takes you to the license plugin that displays a list of your configured license servers. The items on this dashboard are described on page 65 of this Administrator Guide.

## Job Info section

The Job Info section is located in the Grid menu bar.

## Host Job Statistics Viewer

Navigate to the Host Job Statistics Viewer by clicking By Host under the Job Info section of the Grid menu bar. This page shows information about hosts in a cluster.



- **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. Note that you must set how the Down Service Satus should be treated (via Console > THold > Grid Settings tab under Configuration list) for the corresponding option to be available in the Status filter.

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

If graphs have been created for this host, a graph icon appears to the left of the host name. Click the icon to view graphs for the host.

## Host Group Job Statistics

Navigate 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, this page shows information similar to that obtained using the LSF `bhosts` command with condensed host groups.

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

This page shows job information by LSF host group:

- **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 currently set, this field will show N/A. Otherwise, it will show the current value 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. Note that you must set how the Down Service Satus should be treated (via 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 CPU utilization for hosts in this host group.
- **Avg r1m.** The average exponentially-averaged effective CPU run queue length for this host group over the last minute.
- **Avg Effic.** The average efficiency of the host group.
- **Total CPU.** The overall CPU utilization rate of the host group.
- **Max Memory.** The maximum memory consumed 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 used by jobs dispatched to this host group.
- **Run Slots.** The number of job slots used by jobs running on this host group.
- **SSUSP Slots.** The number of job slots used by system suspended jobs on the host group.
- **USUSP Slots.** The number of job slots used by user suspended jobs on the host group.
- **Reserve Slots.** The number of jobs slots used by pending jobs that have job slots reserved within the host group.

You can view Job Statistics graphs by clicking on 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 in neither Unavail, Unreach, Unlicensed, nor 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

Navigate to the Project Viewer by clicking By Project under the Job Info section of the Grid menu bar. This page shows resources in a cluster by project.

The information shown on this 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 used for this project.
- **Pending Slots.** The number of job slots used by pending jobs for this project.
- **Running Slots.** The number of job slots used by running jobs for this project.
- **Avg Effic.** The average efficiency of this project.
- **Max Mem.** The maximum memory consumed by this project.
- **Avg Mem.** The average amount of memory used by this project.
- **Max Swap.** The maximum swap space used by this project.
- **Avg Swap.** The average swap space consumed by this project.

- Total CPU. The overall CPU utilization of this project.

## License Project Viewer

Navigate to the License Project Viewer by clicking **By License Project** under the Job Info section of the Grid menu bar. This page shows resources in a cluster by project.

The information shown on this page is as follows:

- License Project. The name of the license project. Click a licenses project name to show running jobs for this licenses 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 used for this license project.
- Pending Slots. The number of job slots used by pending jobs for this license project.
- Running Slots. The number of job slots used by running jobs for this license project.
- Avg Effic. The average efficiency of this license project.
- Max Mem. The maximum memory consumed by this license project.
- Avg Mem. The average amount of memory used by this license project.
- Max Swap. The maximum swap space used by this license project.
- Avg Swap. The average swap space consumed by this license project.
- Total CPU. The overall CPU utilization of this license project.

## Queue Viewer

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

The information shown on this 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 current 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 used by running jobs in the queue.
- Pend Slots. The number of job slots used by pending jobs in the queue.
- Suspend Slots. The number of jobs slots used by suspended jobs in the queue.
- AVG Pend. The average number of job slots held by pending jobs in the queue.
- MAX Pend. The maximum number of job slots held by pending jobs in the queue.
- AVG Run. The average number of job slots held by running jobs in the queue.
- MAX Run. The maximum number of job slots held by running jobs in the queue.

If you select any queue, you will be directed to a display of all “RUNNING” jobs within that queue.

## View Job Array Listing page

Navigate to the View Job Array Listing page by clicking **By Array** under the Job Info section of the Grid menu bar. This page shows information similar to the LSF `bj obs -A <j ob_i d>` command, but also includes aggregate information for the job array as a whole.

The information shown on this 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 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 used by jobs in the array.
- Avg Swap. The average swap space used by jobs in the array.
- Total CPU Time. The total CPU time used by all started jobs in the job array.

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

## View Job Listing page

Navigate 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 shown on this page is as follows:

- Job ID. The job ID that LSF assigned to the job. Click a job number to view an information page containing 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 current status of the job.
- State Changes. The number of times that the status of the job has 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 has been in the pending state.
- Run. The length of time for which the job has been running.
- SSusp. The length of time the job has been 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

Navigate to the Batch Application page by clicking **By Application** under the **Job Info** section of the **Grid** menu bar. This page allows you to quickly check job details by application.

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

---

### Note:

- LSF Application should be configured first at:  
`$LSF_ENVDIR/l sbatch/<clustername>/confi gdi r/  
l sb. appl i cati ons.`
  - Unused application status are cleared at regular intervals. You can configure this by navigating 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 on the **Application name** link.

You can also click **View Application Graphs** action icon to see the job application graphs. These 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

Navigate 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 via the **Apps** field. Click **Go** when you have completed the filter setting.

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

On the **Submission Details** section of the page, you will see the **Application** details.

---

### Note:

Only users with the **View Application Data** permission can see the **Applications Data**. this 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

Navigate to the Batch Job Groups page by clicking By Group under the Job Info section of the Grid menu bar. This page allows you to quickly check job details by group.

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 on the corresponding Group Name link.

---

### Note:

- If you filter the Group by Level, data from hidden sublevels will be shown in the upper level job group.  
  
The maximum level on the Level dropdown 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 representing statistics on job efficiency, memory stats, pending jobs, running CPU, running jobs, and VM stats. These graphs provide general image of your jobs and lets you control the job flexibly:

- 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 only aggregates job group information by the specified level. If the level is set to 0, the job group data will not be aggregated.

---

### Note:

Only users with the View Job Group Data permission can see the Group Data. This 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 has not been updated you want to have the corresponding graphs removed automatically.
5. Click Save.

## View Job Group in Job Detail page

Navigate 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 via the JGroup field. Click Go when you have completed the filter setting.

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

On the Submission Details section of the page, you will see the Job Group details.

## View Exception Status in Job Detail page

---

### Important:

You need to 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

Once you have configured these parameters, you can submit jobs to the new queue which now includes the exception status definition.

---

Navigate 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 when you have completed the filter setting

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

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

## Enable Pending Reason History Reporting

You can collect and store all pending reasons throughout a job's life cycle. This helps you 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 Platform LSF increases the polling time.
  - For large clusters with over 100K jobs per day, you must configure CONDENSE\_PENDING\_REASON parameter in Platform LSF. Refer to the Platform 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 checkbox.
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 checkbox corresponding to the reason(s) 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:

- General Setting (by RegExp) allows you to enter the reason to be ignored.
- 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, etc.).
- 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, etc.).

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

---

## View Job Pending Reason History page

Navigate 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 once you have completed 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 sorted by the total duration 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 appears in the report when the job state has just 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 located in the Grid menu bar.

### User Viewer page

Navigate to the User Viewer page by clicking Users under the User/Group Info section of the Grid menu bar. This page shows job information pertaining to an LSF user:

- **User Name.** The name of the LSF user. Click a user name to show details of running jobs 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 used by the specified user's jobs.
- **Started Slots.** The number of slots used by jobs submitted by this user and started by LSF. Started jobs can either be running, system-suspended or user-suspended.
- **Pending Slots.** The number of job slots used by the user's pending jobs.
- **Running Slots.** The number of job slots used by the user's running jobs.
- **Effic %.** The average CPU efficiency for jobs submitted by this user.
- **Sys Susp Slots.** The number of job slots used by the user's system-suspended jobs.
- **User Susp Slots.** The number of job slots used by the user's user-suspended jobs.
- **Reserve Slots.** The number of job slots used by the user's pending jobs.

### User Groups Viewer page

Navigate to the User Groups Viewer page by clicking Groups under the User/Group Info section of the Grid menu bar. This page shows job information pertaining to an LSF user group:

- **Group Name.** The name of the LSF user group. Click a user group to show details of running jobs 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 submitted by the specified user group's jobs.
- **Num Slots.** The current number of job slots used by the specified user group's jobs.
- **Started Slots.** The number of slots used by jobs submitted by this user group and started by LSF. Started jobs can either be running, system-suspended or user-suspended.
- **Pending Slots.** The number of job slots used by the user group's pending jobs.
- **Running Slots.** The number of job slots used by the user group's running jobs.
- **SSusp Slots.** The number of job slots used by the user group's system-suspended jobs.
- **USusp Slots.** The number of job slots used by the user group's user-suspended jobs.
- **Reserve Slots.** The number of job slots used by the user group's pending jobs.

## Load Info section

The Load Info section is located in the Grid menu bar.

### Host Load Viewer

Navigate to the Host Load Viewer by clicking Host under the Load Info section of the Grid menu bar. This 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 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 current status of the host.
- **Duration In Current State.** The amount of time that the host spent in its current status.
- **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

Navigate to the Group Load Viewer by clicking Host Group under the Load Info section of the Grid menu bar. This page shows host performance information 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 current 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 on 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 neither Unavail, Unreach, Unlicensed, nor 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 located in the Grid menu bar.

### Closed Admin Host Viewer

Navigate to the Closed Admin Host Viewer by clicking Closed under the Host Info section of the Grid menu bar. This page shows information on 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

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

---

**Note:**

- This page shows the distribution of hosts across different queues in a matrix format.

Note that the very first user who accesses this page may see a blank page. The matrix must be rebuilt to see the details.

- Click **Rebuild** button to rebuild the matrix based on the current configurations.

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

- This page is controlled by realm permissions and is disabled by default.
  - Filters for **Cluster**, **Model**, **Group**, **Type**, **Refresh**, and **Records** are available in this page. Click **Go** to perform the search.
  - Select **HG: Roll-Up** from the Group filter dropdown 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.
- 

This 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

Navigate to the Host Status Viewer by clicking Servers under the Host Info section of the Grid menu bar. This 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 defined for this host, denoted by resource names, and the values of external numeric and string static resources.

## Client Viewer

Navigate to the Client Viewer by clicking Clients under the Host Info section of the Grid menu bar. This page shows information similar to the LSF command `lshosts` except that it only displays LSF clients. In addition to showing current clients that have registered with LSF, it shows all prior clients that have 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. This can be 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.

This page is very helpful in situations where you are leveraging either submission-only clients or floating client configurations.

## Host Group Viewer

Navigate to the Host Group Viewer by clicking Groups under the Host Info section of the Grid menu bar. This 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 belonging 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 located in the Grid menu bar.

### Job History Daily Statistics page

Navigate to the Job History Daily Statistics page by clicking Daily Statistics under the Reports section of the Grid menu bar. This 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 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 used by submitted jobs.
- **User Time.** The amount of user time used by submitted jobs.
- **Start Date.** The start of the reporting period.
- **End Date.** The end of the reporting period.

### Parameters page

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

- Name. Name of the parameter 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)

Navigate 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, as well as the default LSF cluster for which you want to filter data.
- Visual: Control a number of visual settings including the size of the summary icons and the frequency of refresh intervals.

This page contains the following fields:

- Screen Refresh Interval: The interval used by RTM to refresh a page you are currently viewing.
- Exclusion Filter Status: Determines whether or not 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 below are shown in the host dashboard.
- Exclude Host States: Select one or more of the exclusion states to eliminate hosts of that state from appearing in the dashboard. Doing this leaves only the hosts with states you care about in the dashboard. For this to work, you must enable (check) Exclusion Filter Status.

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

- Timespans: Control how your grid summary icons appear. Each user can display thresholds that they believe represent how their grid hosts are behaving.

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

- Default Grid View Timespan: When viewing jobs in the RTM Interface, indicates which default Timespan is in effect. Note that this setting only applies when viewing 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 not shown in the Details page. As there is a large quantity of data collected on LSF jobs, you will likely 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, Job Group, etc.

- 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 collected on LSF jobs, you will likely 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 may also customize the colors that RTM uses when constructing job graphs. The drop-down 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 only shows pending and suspend reasons that have been processed by RTM.

## View 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 located 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 appear 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 on configuring log files in the UNIX host, refer to (specifically, refer to *Enable cluster hosts to appear on the Platform RTM Syslogs page*).

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 top right of the Syslogs page.

This 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, taken from its time stamp.
- **Time.** The time of the log file entry, 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 based on this entry.



# View time series graphs for clusters and hosts

Time series graphs present cluster status and details that depend on time. These are line graphs with time as the X-axis.

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

- View graphs 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 details, see [Viewing modes](#) on page 49.

- View graphs using the action icons.

When 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. For more details on the action icons, see .

- Configure the graph display settings.

Click the Graphs tab, then the Settings tab on the right side of the tabbed menu interface. For more details, see [Settings tab \(from the Graph page\)](#) on page 50.

## Viewing modes

When in the Graphs page, you can view graphs 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 branch (view thumbnails).
- Click a tree branch to view host-specific information in graph form.
- Beside a graph, click the magnifying glass to view more detailed information 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 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)

Navigate to the Graphs > Settings page by clicking the Graph tab, then by clicking the Settings tab on the right side of the tabbed interface. This page allows you to configure the appearance of your graphs and default page settings, and consists of the following sections:

- **General:** Configure general graph options and display formats.
- **Graph thumbnails:** Configure the size of the thumbnails 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

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 you need to take to view non-time series charts for clusters and hosts:

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.

Monitor clusters using Platform RTM

## Working with application license servers

- *Relationship between Platform RTM and license servers* on page 54
- *License Management section* on page 55
- *Configure license pollers* on page 57
- *Add or edit license servers for RTM to monitor* on page 58
- *Monitor license servers* on page 59
- *View time series graphs for license servers* on page 62
- *View non-time series charts for license servers* on page 63

# Relationship between Platform RTM and license servers

Platform RTM is used to monitor and graph LSF resources (including networks, disks, applications, etc.) in a cluster. In graph or table formats, Platform RTM displays resource-related information such as the number of jobs 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 used by Platform Platform RTM for license control. FLEXlm allows licenses to reside on the network instead of a specific host; this allows any registered host in the cluster to access and use a specific licensed application, thereby increasing usage efficiency.

## Start and stop the FLEXlm server

Platform Platform RTM is already configured with the FLEXlm server. If you need to start or stop the FLEXlm server, navigate to `/opt/flexlm/bin`, and then run the following script:

```
./lmgrd -l /opt/flexlm/log/lmgrd.log -c /opt/rtm/etc/rtm.lic
```

Alternatively, copy `/opt/flexlm/bin/lmgrd.init` to `/etc/init.d/lmgrd`, and then use these commands to start and stop the FLEXlm server:

```
service lmgrd start  
service lmgrd stop
```

# License Management section

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

## License Pollers page

Navigate to the RTM License Pollers page by clicking License Pollers under the License Management section of the Console menu bar. This page shows information on 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 checkboxes for the servers on which to perform an action (for example, Delete), then select an action and click Go.

## License Services page

Navigate to the License Services page by clicking License Services under the License Management section of the Console menu bar. This 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 allows you 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. This allows you to search for license services quickly based on the 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 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 current status of the license service. (Status can be Down, Recovering, Up, and Unknown.)
- **Current Time.** The time taken to poll this service.
- **Average Time.** The average time taken to poll this service.
- **Max Time.** The maximum time taken to poll this service.
- **Availability.** The availability of the license service 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, this is N/A.
- **Last Updated.** The last time that the license poller successfully polled the license service. If the license poller never successfully polled the license service, this is N/A.
- **Choose an action.** Select one or more checkboxes for the service on which to perform an action (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. You need the `licensepoller` feature in your Platform RTM license to use license pollers. The following describes the actions you need to take, in chronological order, to configure license pollers.

1. [Set up remote license pollers](#) on page 57
2. [Add or edit license pollers to the Platform RTM Console](#) on page 57

## 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 need to 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_LICENSEPOLLER` guide.

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

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

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 new 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 describing your license poller.
5. Click Create (or Save) to save the settings for your license poller.

## Add or edit license servers for RTM to monitor

You need to 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.

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 new 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 describing 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.
  - 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

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

- [View the status of the license services](#) on page 59
- [View reports on license servers and projects](#) on page 59

## 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.

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.** The clickable icons 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 on 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 current 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 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 on 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.

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 details, refer to [License Inventory](#) on page 60.
- The Feature Usage link shows the licenses in use for each license feature. For more details, refer to [Feature Usage](#) on page 60.
- The Feature Checkouts link shows information on licenses that are currently checked out. For more details, refer to [Feature Checkouts](#) on page 61.
- The Usage Charts link shows the non-time series charts for your licenses. For more details, refer to [View non-time series charts for license servers](#) on page 63.

## FLEXlm Reports section

The FLEXlm Reports section is located in the License menu bar.

### License Inventory

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

- Actions. The clickable icons 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 current status of this product license.

### Feature Usage

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

- Actions. The clickable icons 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

Navigate to the License Checkouts page by clicking Feature Checkouts under the Reports section of the Grid menu bar. This page shows information on 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 current status of the license checkout.
- InUse. The number of licenses checked out.
- Date. The time of the last change in status.

## Create graphs to monitor your license server

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.

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 on Data Query, refer to the Cacti documentation:

<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. These are line graphs with time as the X-axis.

View time series license graphs as you would for any other time series graphs that you use to monitor your clusters. For further details, refer to [View time series graphs for clusters and hosts](#) on page 49.

---

**Note:**

If you are in Tree View or Preview View, remember to click **List View** on the upper right corner of the page before running the procedures presented below.

---

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 checkboxes beside any graph that you want to view.
4. Click View to view all the graphs that you selected.

## View 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, Platform RTM can present this data in graphically-rich formats, such as pie charts and bar graphs.

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

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.





## Working with thresholds and alerts

- [Configure thresholds and alerts](#) on page 66
- [View threshold and alert information](#) on page 68
- [Acknowledge triggered alerts for thresholds](#) on page 70

# 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](#) on page 66
- [Modify threshold settings](#) on page 67
- [Delete thresholds](#) on page 67

## Create a threshold to trigger alerts

You can create thresholds using the following sources:

- [Create a threshold from a graph template](#) on page 66
- [Create a threshold from a host](#) on page 66

### Create a threshold from a graph template

Create a threshold using a graph template as the source.

1. Click the Console tab.
2. Under the Management section of the Console menu bar, click Thresholds.
3. Click Add on the top 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.

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

- 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 your new threshold.

### Create a threshold from a host

Create a threshold using a host as the source.

1. Click the Console tab.
2. Under the Management section of the Console menu bar, click Thresholds.
3. Click Add on the top 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.

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

- 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 your new threshold.

## Modify threshold settings

Make changes to the configuration of thresholds.

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, make desired changes to your threshold configuration.

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

- 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 the your changes to the threshold configuration.

## Delete thresholds

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

1. Click the Console tab.
2. Under the Management section of the Console menu bar, click Thresholds.
3. Click the checkbox 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 configuring thresholds, see [Configure thresholds and alerts](#) on page 66.

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 details, see [Thresholds subtab](#) on page 68.

- Review the status of the hosts being monitored by a threshold.

Click the Thold tab, then the Host Status subtab. For more details, see [Host Status subtab](#) on page 68.

### Thresholds subtab

Navigate to the Thresholds page by clicking the Thold tab, then the Thresholds subtab. This page shows the configured thresholds in your cluster.

This 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 currently 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

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

This page contains the following fields:

- Description. A description of the host. This 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](#) on page 70
- [Acknowledge alerts for multiple thresholds](#) on page 70

## Acknowledge alerts for a single threshold

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.

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

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 checkbox 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.

You can repeat the above 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.

## Metadata Plugin

- [Manage metadata plugin](#) on page 72
- [Define metadata XML configuration file](#) on page 72
- [Define metadata using CSV file](#) on page 73
- [Enable Metadata Settings in Realm Permissions](#) on page 74
- [Import metadata XML or CSV configuration file via RTM Console](#) on page 74
- [View metadata details](#) on page 74

# Manage metadata plugin

The metadata plugin allows you to add metadata for users, user group, host, host group, cluster, job group, queue, project, license-project, application, and queue group.

Concepts to remember:

- LSF Object

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

- Meta Configuration File

An xml file that outlines the attributes of each LSF Object. It is like a blueprint for LSF objects. Only a single meta configuration is permitted.

- Metadata

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

*Example:*

A common use for metadata plugin is when you want to display metadata for host such as geographic location or rack location.

1. Create metadata configuration file.

When defining XML configuration file, refer to the descriptions of metadata configuration file tags:

[Define metadata XML configuration file](#) on page 72

2. Import the XML or CSV metadata file.

Access to the Metadata Settings pages is controlled in Realm Permissions. For details, refer to [Enable Metadata Settings in Realm Permissions](#) on page 74 .

- Import configuration file via 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 via the RTM Console:

[Import metadata XML or CSV configuration file via RTM Console](#) on page 74

## Define metadata XML configuration file

Sample metadata configuration file is located by default at:

```
/opt/cacti/plugins/meta/metadata.conf.xml
```

You can refer to this list of metadata file tags when preparing the configuration file:

- <meta>: this is the enclosing tag for the entire metadata configuration file



- `<object>`: specifies the metadata object type to be configured; the type can be any of the following: 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 popup creates a new tab for each section, allowing you to 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>`: label to be displayed for this column in the UI
- `<description>`: description to be displayed for this column in the UI; this value is displayed when adding a new 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:

- **Text**: commonly used for most columns; the plugin displays the value in the UI
- **URL**: when the data represents a hyperlink, the plugin formats this column as a link
- **Callback**: for this value, the plugin calls the URL provided and displays the result in the UI, without any parsing; this allows dynamic data to be presented as metadata like the current temperature of a host but it does have performance consequences. Note that this tag should 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 should be displayed in the UI, ascending, starting with 1; this value must be unique for the type
- `<summary>`: indicates whether this column should be 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 specified in the metadata UI
- `<popup>`: determines whether this column is displayed in the metadata popup 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 popup.

---

## Define metadata using CSV file

You may choose to load all metadata information using a CSV file. The format of this file must follow this:

- 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 following the order of the header.

*Example:*

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

This is an example of a User object type, where `meta_col1` represents user's full name and `meta_col2` is the user's phone number.

---

**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 popup.

---

## Enable Metadata Settings in Realm Permissions

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

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 on the User Name you want to grant metadata access.
4. Click Realm Permissions tab.
5. Click RTM -> View Metadata and RTM -> Edit Metadata checkboxes to enable these settings.
6. Click Save.

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

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

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 will be loaded.

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

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

- To delete the metadata object(s), select the metadata file(s) on the list then choose Delete action before clicking Go.

## View metadata details

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

## Generic alarms (advanced feature)

- [Configure alarms](#) on page 76
- [Create an alarm using expression](#) on page 76
- [Create an alarm with custom scripts](#) on page 79
- [View alarm log](#) on page 79

## Configure alarms

You must have a good knowledge of the RTM database and structure to be able to use this advanced RTM feature for generating generic alarms. Creating of specific alarms is not supported.

Two types of generic alarms are supported by RTM:

- Expressions

A GUI builder that allows you to create expressions that are to be evaluated for threshold value.

- Scripts

You can create custom scripts that return threshold values. Your script can have very complex queries.

The generic alarm allows you to pull data from any database table and create threshold values for it.

---

### Restriction:

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

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

---

### Important:

Grid Alarms are controlled by realm permission. To set this, Click **User Management** under the Utilities section of the Console menu bar. In the Realm Permissions tab, you can select or deselect the permission type (s) to enable or disable user access accordingly:

- **View Grid Alarms**

Allows user to access the Alarm Details page.

Users can view the alarm details by clicking on the View Alarm Details action icon. Current items that satisfy the conditions of your expressions or scripts are shown on the page.

- **Configure Alarms**

Allows user to create metric, expression, and alarms.

---

## Create an alarm using expression

The expression builder contains two parts: Metrics and Expressions. The GUI builder allows you to create metrics and expression based on a specific table in the database. The ultimate goal of the expression builder is to create the WHERE clause of an SQL statement to be able to retrieve a set of records that is violating the threshold.

The steps below guides you how to create an alarm using expression based on the specific sample problem shown here:

*Problem:* Create an alarm that alerts/emails the Administrator when a host's free disk space has dropped below 100MB.

1. Create an elim in LSF to report free disk space.

This information will be collected by RTM in grid\_hosts\_resources table.

- Build two metrics for `grid_hosts_resources` table. One for the `totalValue` column and the other for the `resource_name` column.

Follow steps in [Create metrics](#) on page 77.

- Create an expression that returns the hosts that will have disk space less than 100MB.

Follow steps in [Create an expression](#) on page 77.

For this example, the `elim` script return the size in bytes so `100MB = 100000000` bytes. SQL Query should be like this:

```
SELECT * from grid_host_resources WHERE resource_name = 'df' AND totalValue < 100000000
```

- Create an alarm by associating it with the newly created expression.

Follow steps in [Create an alarm](#) on page 78.

For this example, set the Grid Alarm Aggregation Method to **COUNT** to see how many hosts are breaching the threshold. You must also set High Value to **0** so that you are alarmed everytime a host breaches this threshold.

## Create metrics

The expression builder starts with creating a set of metrics. The metrics represent a column in the database table.

- Click the Grid tab.
- Under the Grid Alarm section of the Grid menu bar, click Metric.
- In the Grid Alarms Metric page, click Add link from the upper right corner of the page.
- Enter the required field details in the Grid Alarms Metric [edit] page.
  - For Name and Description fields, enter meaningful details for the metric you are creating.
  - For Type field, choose the type of metric you want to collect from the predefined tables for Job, Hosts/Groups, Queue, User/Groups, and Cluster. Optionally, you can select Custom which presents you with a full list of all tables in the database. This determines the available table for selection.
  - For Table field, select the database table where the metric you are creating belongs to. This determines the available columns for selection.
  - For Column field, select the database table column that the metric represents.
- Click Create.

The created metric appears in the Grid Alarms Metric list. You can edit or delete a selected metric.

## Create an expression

Once you have created your set of metrics, you can start building your expression. Each expression can be attached to a single database table just like the metrics. The metrics available to you are only from within that table.

- Click the Grid tab.
- Under the Grid Alarm section of the Grid menu bar, click Expression.
- In Grid Alarms Expressions page, click Add link from the upper right corner of the page.
- Enter the required field details in the Expression [new] page.
- Create your expression.

## Generic alarms (advanced feature)

- For Name and Description fields, enter meaningful details for the expression you are creating.
  - For Type field, choose the type of metric you want to use for this expression.
  - For Table field, select the database table for this expression.
6. Click Create.
  7. In the Expression box, enter your expression. This must conform to the SQL syntax.
  8. Click Check Syntax to check your SQL query.  
Either OK or Bad Syntax result appears. Refer to Cacti log for the actual SQL error.
  9. Click Save.

The created expression appears on the Expression Items list. You can delete or edit the selected expression item. After creating a new expression, you cannot change the database table associated with it.

## Create an alarm

You can start creating an alarm after you have properly defined your expressions or scripts.

1. Click the Grid tab.
2. Under the Grid Alarm section of the Grid menu bar, click Alarm.
3. Click Alarms tab.
4. In the Alarm Management page, select the cluster where you want to apply the alarm.  
This automatically adds a `clusterid` clause to *where* statement.
5. Click Add link from the upper right corner of the page.
6. In the Alarm Settings page, configure the settings based on required details.

Below are some of the filed settings you may want to take note of:

- You should specify the type of data collection (Script or Expression) you want to use for this alarm in the Grid Alarm Type field.
  - In most cases, you should select Alarm Enabled to indicate that the alarm is active and specify the type of aggregation method you want to use in the Grid Alarm Aggregation Method field (such as COUNT, SUM, AVG, MIN or MAX).
  - In the Mandatory Settings section, indicate the trigger for Weekend Exemption and Disable Restoration Email.  
Select Acknowledgement Required if you want the alarm to be acknowledged by the Administrator. If this is checked, the alarm will not return to normal state unless acknowledged by the Administrator.
  - Specify the High Value or Low Value for triggering the alarm and set the Breach Duration for the amount of time the measured value must be in breached before an alarm is raised.
  - If you are using expressions, the details for the Email Subject and Email Body contents are created by RTM. If you are using scripts, you must provide the script for the details.
  - The Notify Accounts let you append or remove items from the two multi select boxes when specifying the accounts to be notified when the alarm is breached. In Extra Alert Emails field, you can also specify extra email accounts that should receive alerts for this alarm.
7. Click Save to create the alarm.

Created alarms appear on the list of alarms. You can acknowledge (if required), edit or delete the selected alarm.

## Create an alarm with custom scripts

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

---

**Note:**

A special variable `|clusterid|` can be used. This is replaced with the selected clusterid at runtime.

---

1. Create the alarm using custom scripts.

Follow steps in [Create an alarm](#) on page 78.

Note that for this method, you must provide these two items when 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 script 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 script and arguments must not exceed 255 characters.

## View alarm log

RTM keeps a logging trail of the alarm activities. Acknowledgement, Triggers, and Clear information details are logged.

1. Click the Grid tab.
2. Under the Grid Alarm section of the Grid menu bar, click Alarm.
3. Click Alarm Log tab.
4. In the Alarm Logs page, you can select the option in the Alarm, State, and Type fields or enter details in Search field for the alarm logs you want to view.
5. Click Go.

Alarm logs matching your search criteria are shown on the list.

Generic alarms (advanced feature)



## Performance and maintenance

# Performance tuning for large clusters

## Store rrd files on a separate disk

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

1. Create a new directory on a separate disk.

For example:

```
cd /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

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

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

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

If you have problems with high disk I/O wait times, enable on-demand rrd file updating. If after enabling this 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.

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

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

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

Partitioning is required if you have a large number of jobs per day or if you want to extend the period of time to keep job summary data.

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

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

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 checkmark.
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 creating a new partitioned table. Increasing the number of partitions means the database will contain more data for overall analysis but also increases the size of the database.

7. Click Save.

## Configure data collection frequency

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
```

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

If you see errors in the `caacti.log` file indicating the LSF APIs are timing out, increase the timeout value.

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 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.

# Database maintenance

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

- [Set up record partitioning in the database](#) on page 85
- [Back up and restore the Platform RTM system settings](#) on page 86

## Set up record partitioning in the database

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. This 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

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 should be removed from the database.

All partitioned data are always retained based on the retention periods you specified under this 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 a large number of jobs per day or if you want to extend the period of time to keep Job Summary data.
- **Partition Size**—Specify a period of time between different partitions.

Increasing the period means the database will contain 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.

The partitioned tables have a numerical suffix (`_v000`, `_v001`, ..., `_vnnn`) with the newest table having 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 perform operational backups of the Cacti database so that you can restore the Platform RTM system at a later date. This 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 will restore the Platform RTM system and graph settings and revert to your original, backed-up settings.

---

### Restriction:

The Platform RTM Console will not restore all of the job performance data of the clusters that Platform RTM monitors — any such data still in Platform RTM will likely be out of date.

If you want to perform a full backup of the entire database, you will need to use the `mysql dump` command. In addition, you should not use the MyISAM table structure if you use the `mysql dump` command on a regular basis. For more details, refer to MySQL documentation (for example, refer to the MySQL web site 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 using the Platform RTM Console:

- [Back up the Platform RTM system using the Platform RTM Console](#) on page 87
- [Enable automatic data archiving](#) on page 88
- [Restore the Platform RTM system using the Platform RTM Console](#) on page 88

The following files are backed up:

- `rtm.lic`: Platform RTM license file. This file is not restored automatically when you restore your configuration within the Platform RTM Console.
- `lsfpollerd.conf`: Database file containing the credentials
- `lsf.conf`: The `lsf.conf` file associated with each cluster.
- `ego.conf` (for LSF 7.x and LSF 8.0 clusters only): The `ego.conf` file 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 are backed up except for the following:

- `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

Prior to upgrading to a newer version of Platform RTM or to a fully licensed version, you must first back up the existing Cacti database.

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

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 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 you want the backup to take place
  - Backup Generations—Specify the number of backup file copies you want to maintain
  - Database Backup Location—Provide a location if a backup directory does not yet exist. (If the directory is found/exists, the message “[OK: DIR FOUND]” displays under the directory field.)
  - Post Backup Command—Provide the post backup script to be run once 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. Once downloaded, verify that the .tgz file contains the following files:
  - 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)

Once 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

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.

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 prior to 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

After you have completed any scheduled server maintenance, or finished upgrading your Platform RTM version or license, 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.

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 using the Platform RTM Console. However, previous versions of Platform RTM require that you manually backup your database. The following steps describe the manual backup and upgrade process:

1. [Back up the web site](#) on page 89
2. [Back up the MySQL database](#) on page 89
3. [Back up the Platform RTM poller](#) on page 90
4. [Restore the Platform RTM license file](#) on page 90
5. [Upgrade the Cacti web site and grid plugin](#) on page 90
6. [Upgrade the MySQL database](#) on page 91
7. [Upgrade the Platform RTM poller](#) on page 91



## Back up the web site

This is recommended whenever you upgrade Cacti. In order to do this, you should backup all directories except the `log` and `rra` directories. If you are not concerned about the size of the backup, you can continue to backup those directories though. Below 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 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 some very important Cacti and Platform RTM tables to disk. It does not back up some rather large tables in Platform RTM simply due to their size and purpose. The Cacti/Platform RTM tables *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 only backs up the file path to the key file.

Even though the contents of the tables are not backed up, their structure is. In order 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). Once you have verified that, open a shell and perform the following steps:

```
cd /var/www/html/cacti/plugins/grid
```

```
php -q database_backup.php
```

Running that utility will create 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, in order to restore full functionality to the system.

## Back up the Platform RTM poller

In order to support the ever growing list of features and to provide for bug fixes, the Platform RTM pollers will have to be upgraded from time to time. Therefore, prior to upgrading Platform RTM, you should first make a backup copy of your Platform RTM pollers.

In order to backup your pollers, execute the following commands on each poller, or simply backup 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
```

## Restore the Platform RTM license file

The Platform RTM license file is automatically backed up as part of the database backup procedure, but it is not restored during the database restore procedure. To manually restore a license file, extract the license file from the backup .tgz archive file. For example, if you installed Platform RTM to /opt/rtm, open a shell and perform the following steps:

```
tar zxvf cacti_db_backup_num.tgz cacti_backup/rtm/etc/rtm.lic
cp cacti_backup/rtm/etc/rtm.lic /opt/rtm/etc/rtm.lic
service lmgrd restart
```

## Upgrade the Cacti web site and grid plugin

In order to upgrade the Cacti web site, you need to 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 plugins.
2. You must know whether or not you have modified the default Platform RTM poller scripts. If you have done so, you will need to make a separate backup of your <path\_to\_cacti>/scripts directory. If you have added new scripts that are not a part of the standard Platform RTM deployment, you do not need to back up this directory.
3. If you have modified default Platform RTM Data Queries, you need to backup your <path\_to\_cacti>/resource directories. Again, if you have not modified the default installation of Platform RTM Data Queries, you do not need to proceed with this step.

Now that you have done this, note the ownership of all files in Platform RTM. The Platform RTM installer recommends that the Apache site be owned by the cacti user account, and that the Apache binaries be run by the same. If this is the case, then all files in the /var/www/html/cacti directory and sub-directories should be "cacti:cacti". Regardless of file ownership, not the ownership and permissions.

When you receive your Platform RTM upgrade, there should be two directories in the upgrade path. They include both the "cacti" and "poller" directories. In order to upgrade the Cacti Web Site and Grid Plugin, simply run the following command:

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

Then verify that your config.php remains unchanged. Then refresh your browser that should already be pointed to the Platform RTM web site. If the web site does not come up, contact Platform Technical Support ASAP.

## Upgrade the MySQL database

Once you have upgraded your Cacti Web Site and Grid Plugin, 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_rusage` do not typically undergo structural changes. If one of these larger tables were to require modification, you would be informed up front in order to advise you of the downtime that may be required to perform any upgrade to those tables.

---

In order to upgrade your database schema, run the following commands from a shell:

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

If, after running the database upgrade script, you receive error messages, contact Platform Technical Support ASAP.

## Upgrade the Platform RTM poller

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

They should typically be owned by “cacti:cacti”. If you have chosen another method, please note those file permissions.

2. Execute the following commands to upgrade the pollers:

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

## Set up a remote database host

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 Platform RTM 8.0 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.

By default, the Platform RTM host runs as the database host. You may 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:

You should perform these steps during the initial setup of your Platform RTM installation. If you choose to perform these steps after Platform RTM has been in operation for a while, the database may take a significant period of time to transfer to another host, during which the database will be down.

For more details on effectively transferring the database to another host under these conditions, contact Platform Computing.

---

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

Follow the steps described in [Database maintenance](#) on page 85 (specifically, refer to *Back up and restore the Platform RTM system settings > Back up the Platform RTM system using the Platform RTM Console*).

2. Prepare the Platform RTM host for the database transition.
  - a) Log into 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 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 will 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, navigate to the line where the database host name is being declared and change "local host" 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/lfsfversion\_number/bin/grid.conf
  - /opt/rtm/etc/lfspollerd.conf
  - /etc/rsyslog.conf
  - /opt/cacti/plugins/syslog/config.php
3. Set up the new database host for Platform RTM.
    - a) Log into 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 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 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 into the Platform RTM host.
    - b) Restart the essential services by running the following commands:
 

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

```
service lsfpollerd status
```

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

## Issues to consider

Platform RTM has some issues that you should keep in mind.

### LSF ports that RTM requires

RTM requires access to the LIM (load information manager) port. If you do not specify this port, Platform RTM will not be able to communicate with the LSF cluster.

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

RTM does not need access to the `sbatchd` (slave batch daemon), `mbatchd` (master batch daemon), and `RES` (remote execution server) ports, as it does not need to communicate with these LSF components.

### Known issues

For a list of the latest known issues, refer to the *Release Notes for Platform RTM*.

### Troubleshooting and FAQs

For a list of troubleshooting tips and frequently asked questions, refer to the RTM site: <http://my.platform.com/products/platform-lsf-family/platform-rtm/faq>