
Platform RTM Administrator Guide for SAS

Platform RTM
Version 2.0.7
Release date: December 2011



Copyright

© 1994-2011 Platform Computing Corporation.

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

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

We’d like to hear from you

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

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

Document redistribution and translation

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

Internal redistribution

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

Trademarks

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

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

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

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

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

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

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

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

Third-party license agreements

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

Contents

1	About Platform RTM	5
	Introduction to Platform RTM	6
	About the user interface	7
2	Getting started	13
	Activate Platform RTM	14
	Enable Grid Control on an LSF Cluster	16
	Configure date, time, and license information	18
3	Using Platform RTM to monitor the cluster	19
	View LSF cluster and job information	20
	View threshold and alert information	38
	View UNIX log file entries	40
	View time series graphs for clusters and hosts	41
	Monitor high availability applications	43
4	Configuring Platform RTM	45
	Configure cluster management	46
	Configure cluster interaction	52
	Configure date, time, and license information	57
	Configure thresholds and alerts	58
5	Administering Platform LSF	61
	Add or edit LSF clusters for Platform RTM to monitor	62
	Control an LSF cluster	65
	Configure an LSF cluster	67
	LSF log files	72
	High availability applications	74
6	Performance and maintenance	81
	Database maintenance	82
	Issues to consider	85

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 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.

Relationship between LSF and Platform RTM

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, or job owner), or the hosts on which the jobs ran.

About the user interface

For the most part, Platform RTM follows the design cues from the original Cacti product. This section describes the details common to all elements of the Cacti user interface, allowing you to more easily navigate its functionality.

- [Tabbed interface](#) on page 7
- [Page navigation and layout](#) on page 8
- [Selection filtering](#) on page 9
- [Action icons](#) on page 10

Tabbed interface

There is a tab for each major area of functionality within the product.

The following example shows the tabs and interface display when you click the Graphs tab, which displays all the tabs.



The following table describes components of Cacti's default user interface configured to include Platform RTM.

Interface component	Description
Tab: Config	Opens the Config page. Access Cacti and Platform RTM administration functions including graph creation and management, templates, grid settings, date and time settings, Platform RTM license updates, and utilities.
Tab: Graphs	Opens the Graph page. View graphs to which your Platform RTM Administrator has given you access.
Tab: Alerts	Opens the Thresholds page. View information about the configured alerts and thresholds in your cluster.
Tab: Grid	Opens the Grid page. View information about your LSF cluster, hosts, queues, and submitted jobs. You can also control your LSF cluster here.
Tab: Syslogs	Opens the Syslogs page. View entries from the UNIX log files located in the <code>/var/log</code> directory in each host in the clusters that Platform RTM monitors. This tab only exists on Linux versions of Platform RTM.
Tab: HA	Opens the High Availability Application dashboard. View high availability (HA) applications and monitor application availability.

Interface component	Description
Tab: Settings	<p>Allows you to customize either the layout of your graphs or the on-screen presentation of your grid.</p> <p>At times there are tabbed options visible to the right of the Settings tab. These change depending upon which page is opened. For example, if you are on the “Graphs” page, tab options are available to allow you to switch between graph views (tree, list, or preview).</p> <p>You may not see the Settings tab at all if the Cacti administrator has restricted your access.</p>
Navigation bar	<p>This is the area just below the tabs where you can find navigational “breadcrumbs”. The circular button on the left allows you to hide and show the grid menu bar (described in the next section). The area to its right allows you to easily navigate up a menu level when you are inside of a menu item.</p>
Logout link	<p>This link is to the right of the navigation bar. Click this link to log out of the system.</p>

Page navigation and layout

As a graphically-rich application, Platform RTM presents several user interface elements that allow for easy page navigation and layout.

- [Menu bars](#) on page 8
- [Page navigation](#) on page 8
- [Headers and sorting](#) on page 9
- [Page layout preferences and customizations](#) on page 9

Menu bars


Menu bars run vertically along the left sides of the Config and Grid tabs. Use the Config menu bar and the Grid menu bar to respectively access administration tools and functions, or to view information about your LSF Cluster and submitted jobs.

Note that your Platform RTM administrator may hide or show various menu items; you may not have access to all of the areas described.

To hide or show various menu items, click User Management in the Cacti Utilities section of the Config menu bar. After clicking the user account to modify, the Realm Permissions subtab selects the menu items that the user can access. For more details, see the Cacti documentation (<http://cacti.net/documentation.php>),

Page navigation

Using the navigation bar, you can move from page to page within a display area.



Option descriptions:

- Click <<Previous to return to a previous page in the list.
- Click Next>> to move forward a page.
- Click any page number in the center of the bar to immediately go to that page.

Headers and sorting

Actions	Host Name	Cluster	Type	Model	Load/Batch	CPU Fact	CPU Pct	RunQ 1m	Max Slots	Num Slots	Run Slots**	SSUSP Slots	USUSP Slots	Reserve Slots
---------	-----------	---------	------	-------	------------	----------	---------	---------	-----------	-----------	-------------	-------------	-------------	---------------

Click a column heading to sort the contents of the display area based on your selection.

The default sort order is controlled at the system level and is biased towards the most likely sort order for that information. Clicking on a column heading reverses the sort order. Some columns may not appear sortable—this is the normal behavior

Page layout preferences and customizations

Depending on the configurations made by the Platform RTM Administrator, you may or may not be allowed to modify personal grid settings. These settings control a user's default environment and interface display. From the Settings page, you can control graph colors, change the default number of rows to display in any section, or show/hide certain fields within the interface.

Click the Settings tab, located on the right side of the tabbed interface, to display the Settings page. This tab is accessible after you click the Graphs, Alerts, Grid, or Syslogs tabs:

- If you click the Settings tab after clicking the Graphs, Alerts, or Syslogs tabs, the Graphs > Settings page displays.

This page allows you to configure Platform RTM graph display settings. For more details, refer to *Settings tab (from the Graph page)* in [View time series graphs for clusters and hosts](#) on page 41.

- If you click the Settings tab after clicking the Grid tab, the Grid > Grid User Settings page displays. For more details, refer to *Settings tab (from the Grid page)* in [View LSF cluster and job information](#) on page 20.

This page provides access to numerous other tabbed pages where you can configure, in detail, your Grid tab display preferences.

Selection filtering

Within various menus, you can filter information that you want displayed. For example, you might filter by cluster, user, status, etc.

The screenshot shows a 'Batch Job Filters' panel with the following fields and values:

Cluster:	All	User:	All	UGroup:	All	Status:	RUNNING	Effic:	All
Queue:	All	Host:	All	HGroup:	All	Records:	30	Except:	N/A
JobID:		Search:		<input checked="" type="checkbox"/> Dynamic	go	clear	export		

Caution:

Platform RTM now uses unbound memory and execution time for filtering job details using the All status filter. It may occupy a lot of memory if the number of job records is huge. Error appears when system runs out of memory.

Click the inverted green triangle along the title bar to hide or show the selection filter. When hidden, your viewable area increases, but filter options are not lost. Once filter options are set, the displayed information updates to only include the selected items.

Selection filters operate using the “AND” operator. In the example above, if you select the Status RUNNING and the User John, only John's running jobs show in the display area.


Button descriptions:

- Click Go to refresh the page using the current filter criteria.
- Click Clear to return the filters to their default values.
- Certain filters also include an Export button. Use this to export as many as ten thousand records into CSV format using your filter criteria.

The following is the information that this filter displays:

- Warning and Alarm Efficiency. Efficiency is a measure of how well an application utilizes its stated CPU request. It is calculated by dividing the actual number of CPUs used by the requested number of CPUs.
This measure requires the application to be properly integrated with LSF to report this data.
- Flapping is a measure of job state changes. If a job changes state too often, this may indicate a problem in the pre-execution or the last execution host to which the job was submitted. Optimally, the job will change state three times: PEND, RUNNING, FINISHED.
- Job dependencies.
- Invalid job dependencies.
- Exited jobs
- Exclusive jobs
- Interactive jobs

Search field



The Search field allows you to specify a free-format search string. The search only looks for key data fields that cannot easily be found using the drop-down filters.

Use this field to select key fields in the data you are currently viewing. For example, on the Job Info > By Host page, use the Search field to filter the Host Name field. On the Job details page, use the Search field criteria to filter the job ID and name.

Time span selection bar
































On certain pages you can select a time span to view graphs and completed job details for a selected time range (for example, Grid > Job Info > By Array). The Presets field allows the selection of data between common time intervals, such as the last day, hour, week etc. Calendar links (beside to To and From fields) let you define custom time/date ranges.

The arrows on the right side of the Time Span Selection Bar allow you to either advance or go back an amount of time specified within the corresponding drop-down list. For example, if you are currently looking at jobs that finished in the last day, and you click the left arrow on the time shifter, the jobs that completed during the previous 24 hours now display. Alternatively, if you select 1 Week from the drop-down list and then click the left arrow, jobs over a 24-hour period from the previous week display.

Action icons

Many pages have an Actions column in the header. Under this column are various icons indicating the type of action available from within this page. Some common action icons are described in the following table.

Icon	Description	Icon	Description
	View queues		View host job detail
	View users		View graphs
	View batch hosts		Zoom into graph
	View batch host groups		Display jobs in range
	View active jobs		View graph or source properties
	Export data to CSV format (open or save to file)		Return to top of page
	View license checkouts		Connect to remote host (opens a console window)
	Create threshold		Edit threshold
	Disable threshold		Enable threshold
	Acknowledge threshold (stop triggering alerts)		Reset acknowledged threshold (resume triggering alerts)
	Toggle threshold rules		Add a syslog alert rule
	Add a syslog removal rule		Copy high availability application
	Create a time-based configuration		View host logs (displays the available logs in the host).
	Head host logs (displays the earliest entries in a log file).		Tail host logs (displays the latest entries in a log file).
	Download host logs.		

Getting started

This chapter describes the steps you need to take to set up Platform RTM to function properly after installation.

For more information on installing Platform RTM, refer to the *README for Installing Platform RTM for SAS on Windows* or the *README for Installing RTM for SAS on Linux*, depending on the operating system of the intended Platform RTM host.

Activate Platform RTM

Prerequisites:

- You have a valid license key for Platform RTM.
- Your local host (that is, the computer that you will use to access the Platform RTM Console) is running a supported web browser.

Platform RTM supports Firefox 3.6 or later and Internet Explorer (IE) 8 or later.

The following describes the initial post-installation and configuration tasks that you need to perform for Platform RTM to function properly.

1. Log into your local host.
2. Add licenses to your Platform RTM installation.
 - a) Use your web browser to launch the Platform RTM Console.

The URL for the Platform RTM Console is **http://rtm_host/rtm** where *rtm_host* is either the Platform RTM host name or IP address.

For example, if your Platform RTM host name is `hostA.example.com` with IP address `192.168.1.5`, you can use either of the following URLs to launch the Platform RTM Console:

- **http://hostA.example.com/rtm**
- **http://192.168.1.5/rtm**

- b) Specify the Platform RTM administrator name and password.

The default administrator name is **admin**, and the default password for this account is **admin**.

To see the License Info page, click the Admin tab, then the License subtab.

- c) Add the license file to Platform RTM.

You can add the license file to Platform RTM using one of the following methods:

- Click Browse and navigate to your license file.
- Open your license file with a text editor and copy the text to the clipboard, then paste this text to the License Text field in the License Info page.

- d) Click Save to save the license file to Platform RTM.

3. Add LSF clusters for RTM to monitor.

Follow the steps described in `Add or edit LSF clusters for Platform RTM to monitor` in the Platform RTM Administration guide for every LSF cluster that you want Platform RTM to monitor.

4. Add license servers for Platform RTM to monitor.

Follow the steps described in `for every license server that you want Platform RTM to monitor`.

5. Optional. Set up the Platform RTM database on a remote host.

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.

Follow the steps in the Platform RTM Administration guide to set up the Platform RTM database on a remote host.

6. Optional. Click the Config tab to verify and configure the date and time of your Platform RTM host.

You can also update or upgrade your Platform RTM license using this tab. For more information on the Config tab, refer to .

Enable Grid Control on an LSF Cluster

Enable grid control on an LSF cluster to grant Platform RTM access to that cluster. This allows you to run LSF commands on that cluster.

Prerequisites: The LSF cluster and the RTM host must meet the following requirements:

If the Platform RTM host is a Linux host:

- The LSF master host is a Linux, AIX, HPUX, or Solaris host with `sh` or `bash` installed.
- The Platform RTM host is a LSF server, or is added to the LSF cluster as an LSF client.
- The RTM host has `rsh` or `ssh` access to the LSF master host.
- The LSF master host uses at least one of the following methods of authentication and meets the corresponding requirements:
 - `ssh` password authentication: You are asked for the password of the LSF master host root user each time you invoke a cluster control action.
 - `ssh` private key authentication
 - You created an `ssh` public key pair by running `sshkeygen -t rsa` on the RTM host as root, then adding the public key to the `authorized_keys` file of the LSF master host root user.
 - The LSF master host has password-less authentication (`ssh` private key authorization or `rsh`) available with all other hosts in the LSF cluster.
 - `rsh` password-less authentication
 - The `.rhosts` file in the LSF master host specifies the root user of the RTM host.
 - The LSF master host and the RTM host both have the incoming TCP port 514 open.
 - The LSF master host has password-less authentication (`ssh` private key authorization or `rsh`) available with all other hosts in the LSF cluster.

If the Platform RTM host is a Windows host:

- The LSF master host is a Windows host.
- The Platform RTM host is a LSF server, or is added to the LSF cluster as an LSF client.
- The RTM host has WinRS installed and configured correctly.

For details on installing Platform RTM with WinRS, refer to the *README for Installing RTM for SAS on Windows*.

- The LSF master host has WinRM 2.0 installed to allow remote execution.

For details on configuring WinRM to allow remote execution, refer to the *README for Installing RTM for SAS on Windows*.

The following describes the steps that you need to perform to enable grid control and allow you to control LSF clusters using the Platform RTM Console.

1. Click the Config tab.
2. Enable grid control for each applicable user in the RTM host.
 - a) Under the Utilities section of the Config menu bar, click User Management.
 - b) Click the name of the user for which you want to enable grid control.
 - c) In the Realm Permissions section, select the Cluster Control Management field, if it is currently unchecked.
3. Under the Grid Management section of the Config menu bar, click Clusters.

4. Enable grid control on LSF clusters.

Perform the following for each cluster that you want to control:

- a) Click the name of the cluster that you want to control.

The Cluster Edit page displays.

- b) Click Control tab.

The User Authentication settings section appears.

- c) In the User Authentication settings section, specify the settings for the Primary LSF administrator account in the LSF master host.

To ensure that RTM has access to the appropriate LSF commands, you must consider the following:

- The specified Primary LSF administrator user name is the name of the LSF administrator account in the LSF cluster for which you are enabling grid control. You must specify the username of the Primary LSF Administrator for the LSF machine. You need to set the cluster username before executing cluster via advocate, otherwise, invalid credentials/no username specified error appears.

This account is used by the Host-, Queue-, and Job-level controls using `eauth` in the LSF master host to invoke the control actions. After saving these settings, this user name is created as a disabled Unix local account in the RTM host.

- If you are connecting to the LSF master host using `ssh` private key authentication, you need to provide the private key path pointing to the private key file. As shown in the prerequisites, the public key of this file is added to the `authorized_keys` file of the LSF master host root user.
 - The LSF server top directory is the top-level LSF installation directory (`LSF_TOP`).
- d) Click Save to apply your changes.

Configure date, time, and license information

Click the Config tab and navigate to the Cacti Configuration section of the Config menu bar on the left side to configure the date, time, and license information.

The following is a list of actions that you can do to configure date, time, and license information:

Timezone subtab

Note:

This tab does not exist on Windows versions of Platform RTM.

Navigate to the Cacti Settings (Timezone) page by clicking Cacti Settings in the Cacti Configuration section of the Config menu bar, then by clicking the Timezone tab. This page defines the time zone, and current date and time. You can also specify an NTP server.

Note:

If the server is not able to synchronize the date/time, you can manually set the date/time. NTP overrides any manual settings once the server is able to synchronize.

This page contains the following fields:

- Timezone Setting: Set this to your local time zone.
- Date Setting: Select the current date.
- Time Setting: Enter the current time.
- NTP Server: Specify a preferred NTP server.

After changing any of these settings, RTM restarts the system services.

License Info page

Navigate to the License Info page by clicking Licensing in the Cacti Configuration section of the Config menu bar. The first time you log on to the RTM Console, you must provide licensing information from this page. Use this page if your license expires and you need to update it, or if you wish to upgrade your demo license to a full-feature version.

You can either browse to the location of your license file, or you can copy and paste the text from your license file into the appropriate field on this page. Click Save to complete the license update.

If you use SAS licenses for all of your LSF clusters, you can use a SAS license for Platform RTM. If at least one of your LSF clusters uses a Platform licenses, you must obtain and use a Platform license for Platform RTM.

Using Platform RTM to monitor the cluster

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 52.

For more information on configuring LSF clusters using Platform RTM, see [Add or edit LSF clusters for Platform RTM to monitor](#) on page 62.

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 20.

- 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 Cacti Utilities section of the Config 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 22.

- 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 31.

- 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 32.

- 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 33.

- Review daily statistics 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 35.

- 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 36.

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 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"> • ok
	<ul style="list-style-type: none"> • locked • locked-user • locked-window • locked-master






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

Table 2: Batch Status

Icon	Status
	<ul style="list-style-type: none">• ok• Any closed except for admin
	<ul style="list-style-type: none">• closed-admin
	<ul style="list-style-type: none">• unavail• unlicensed• unreach

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.
- **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 Config 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:
`SLSF_ENVDIR/lsbatch/<clustername>/configdir/lsb.appl icat ions.`
 - 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.** On Unix, the idle time of the host, in minutes. On Windows, the time a screen saver has been active on the host.
- **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.** On Unix, the idle time of the host, in minutes. On Windows, the time a screen saver has been active on the host. The average idle time is computed for all hosts with the group.
- **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, Chkpt/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.

-

View threshold and alert information

Click the Alerts tab to view threshold and alert information.

For more information configuring thresholds, see [Configure thresholds and alerts](#) on page 58.

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 Alerts tab, then the Thresholds subtab. For more details, see [Thresholds subtab](#) on page 38.

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

Click the Alerts tab, then the Host Status subtab. For more details, see [Host Status subtab](#) on page 38.

Thresholds subtab

Navigate to the Thresholds page by clicking the Alerts 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 Alerts 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.

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 41.

- 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 [Action icons](#) on page 10.

- 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 42.

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 Config tab. Devices can be added to the tree and modified from the Config 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.

Monitor high availability applications

Click the HA (High Availability) tab to view high availability application information and to monitor the availability of these applications.

The HA tab has one section, the Application dashboard, which provides an overview of high availability applications that run on clusters that Platform RTM monitors.

For more information on high availability applications, refer to [High availability applications](#) on page 74.

Applications dashboard

Navigate to the Applications dashboard by clicking the HA tab.

The Applications dashboard displays useful information about the status of your high availability applications. By default, this page displays up to 30 high availability applications and automatically refreshes every 60 seconds.

The Applications dashboard shows the following information:

- **Process ID:** The application process ID (PID), if there is one assigned. Applications are assigned a PID once it has started.
- **Application Name:** The name of the high availability application.
- **Application Version:** The version of the high availability application.
- **Cluster:** The cluster to which the high availability application belongs.
- **Status:** The current status of the high availability application process. For further details, refer to .
- **Primary Host:** The host on which the high availability application will run.
- **Failover Host:** The host on which the high availability application will run if the primary host is down. Once the application has run on the failover host, the application will continue to run on the failover host even if the primary host is up again.
- **Running Host:** The host in which the high availability application is currently running.
- **Start Time:** The time in which the high availability application started. The default format shows the time stamp in which the application started, but you can also configure Platform RTM to display the amount of time that the application has been running.
- **Choose an action:** Select one or more checkboxes for the high availability applications on which to perform an action, then select an action and click Go.
 - **Restart** - This will restart the high availability application on the original host if it is up.
 - **Migrate** - This will start the high availability application on the other host (the failover host) if it is up. It takes a while to migrate an application. It is recommended to migrate only a few applications at a time to avoid timeout.
 - **Start** - This will start the high availability application based on the application definition.
 - **Stop** - This will stop the high availability application.

Using Platform RTM to monitor the cluster

Configuring Platform RTM

Configure cluster management

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

The Config menu bar has the following sections: Create (graphs), Management, Grid Management, Collection Methods, Templates, Import/Export, Cacti Configuration, and Cacti 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 46.
- 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 49.
- Grid Management: Pollers, Clusters, LSF, HA, 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 Config menu bar.

Thresholds page

Navigate to the Thresholds page by clicking Thresholds under the Management section of the Config 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_fiel d_name	<custom_custom_fiel d_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.
Threshold trigger time	Not available	<TIME>	The time in which the threshold triggered this alert.

Placeholder name	Tag for threshold name	Tag for alert email template	Description
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	<TRI GGER>	The threshold trigger value.
Threshold graph ID	Not available	<GRAPHI D>	The ID of the threshold graph.
Threshold duration	Not available	<DURATI ON>	The duration of the threshold.
Threshold details URL	Not available	<DETAI LS_URL>	A URL to the threshold details page, which is a list of hosts that breached this threshold.
Threshold breached items	Not available	<BREACHED_I TEMS>	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 Config menu bar.

Pollers page

Navigate to the Pollers page by clicking Pollers under the Grid Management section of the Config 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/lsfversion/bin` for Linux or `C:\Program Files\Platform Computing\RTM\lsfversion\bin` for 32-bit Windows). 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 Config 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 62.

- **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 Config > 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 Config > 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.

Grid LSF Configuration page

Navigate to the Grid LSF Configuration page by clicking LSF under the Grid Management section of the Config menu bar. This page shows LSF cluster configurations that you created using the Platform RTM Console.

For more details on configuring an LSF cluster, refer to [Configure an LSF cluster](#) on page 67.

HA Configuration page

Navigate to the HA Configuration page by clicking HA under the Grid Management section of the Config menu bar. This page shows high availability (HA) applications available on LSF clusters that Platform RTM monitors.

For more details on configuring high availability applications, refer to [High availability applications](#) on page 74.

Utilities page

Navigate to the Utilities page by clicking Grid Utilities under the Utilities section of the Config 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 82 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 Config tab and refer to the Config menu bar on the left side to configure cluster interaction.

The Config menu bar has eight sections: Create (graphs), Management, Grid Management, Collection Methods, Templates, Import/Export, Cacti Configuration, and Cacti 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 Config menu bar. For more details on the subtabs in this page, see [Grid Settings page](#) on page 52.
- Configure how Platform RTM interacts with the clusters.
Click Cacti Settings under the Configuration section of the Config menu bar. For more details on the subtabs in this page, see [Cacti Settings page](#) on page 55.

Grid Settings page

Navigate to the Grid Settings page by clicking Grid Settings under the Cacti Configuration section of the Config 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.

- **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.
- Alerts: 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 Cacti Settings under the Cacti Configuration section of the Config 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.

Note:

This option is not required for Windows hosts.

- **Misc:** Cacti Log Viewer for User

This subtab only appears on Linux hosts.

- **Alerting:** 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 `tel net` 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.

Configure date, time, and license information

Click the Config tab and navigate to the Cacti Configuration section of the Config menu bar on the left side to configure the date, time, and license information.

The following is a list of actions that you can do to configure date, time, and license information:

Timezone subtab

Note:

This tab does not exist on Windows versions of Platform RTM.

Navigate to the Cacti Settings (Timezone) page by clicking Cacti Settings in the Cacti Configuration section of the Config menu bar, then by clicking the Timezone tab. This page defines the time zone, and current date and time. You can also specify an NTP server.

Note:

If the server is not able to synchronize the date/time, you can manually set the date/time. NTP overrides any manual settings once the server is able to synchronize.

This page contains the following fields:

- Timezone Setting: Set this to your local time zone.
- Date Setting: Select the current date.
- Time Setting: Enter the current time.
- NTP Server: Specify a preferred NTP server.

After changing any of these settings, RTM restarts the system services.

License Info page

Navigate to the License Info page by clicking Licensing in the Cacti Configuration section of the Config menu bar. The first time you log on to the RTM Console, you must provide licensing information from this page. Use this page if your license expires and you need to update it, or if you wish to upgrade your demo license to a full-feature version.

You can either browse to the location of your license file, or you can copy and paste the text from your license file into the appropriate field on this page. Click Save to complete the license update.

If you use SAS licenses for all of your LSF clusters, you can use a SAS license for Platform RTM. If at least one of your LSF clusters uses a Platform licenses, you must obtain and use a Platform license for Platform RTM.

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 58
- [Modify threshold settings](#) on page 59
- [Delete thresholds](#) on page 59

Create a threshold to trigger alerts

You can create thresholds using the following sources:

- [Create a threshold from a graph template](#) on page 58
- [Create a threshold from a host](#) on page 58

Create a threshold from a graph template

Create a threshold using a graph template as the source.

1. Click the Config tab.
2. Under the Management section of the Config 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 Config tab.
2. Under the Management section of the Config 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 Config tab.
2. Under the Management section of the Config 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 Config tab.
2. Under the Management section of the Config 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.

5

Administering Platform LSF

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 62
2. [Add the RTM host to the LSF cluster as an LSF client](#) on page 63

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 62
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 62
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 Config tab.
2. Under the Grid Management section of the Config 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.

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 63.

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 install/[ion] directory.

For example,

Linux: `cd RTM_INSTALL_PATH/cacti/plugins/grid`

Windows (64-bit): `RTM_INSTALL_PATH\Cacti\plugins\grid`

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 6.2
- 2: LSF 7.0.2
- 3: LSF 7.0.5
- 4: LSF 7.0.6

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=4 --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 63.
2. Verify that the new cluster is added to RTM using the RTM Console by clicking Clusters in the Config menu bar and checking that the new cluster is up.

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

Since you can enable the Platform RTM installer to automatically add the host to the LSF cluster automatically, you can normally skip this task.

However, if you did not have the Platform RTM installer add the host to the LSF cluster, and the host is not an LSF server or a client, you need to manually add the Platform RTM host to the cluster as an LSF client to give Platform RTM access to LSF cluster data. You also need to do this to any LSF cluster that you did not use the Platform RTM installer to add.

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 `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.

- Linux: Edit the `/etc/hosts` file.
 - Windows: Edit the `C:\WINDOWS\system32\drivers\etc\hosts` file.
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.

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)
- `lsadmi n 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) - on Linux hosts only
- 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) - on Linux hosts only

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.
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 20.

Configure an LSF cluster

Platform RTM allows you to configure LSF clusters, hosts, queues, and resources by using the Platform RTM Console to edit certain parameters in the corresponding LSF configuration files. While it does not support every LSF configuration parameter, Platform RTM allows you to perform the most common configuration tasks.

Navigate to the Grid LSF Configuration page by clicking LSF under the Grid Management section of the Config menu bar. This page allows you to configure LSF clusters, hosts, queues, and resources.

The following is a list of actions you can take to configure the LSF cluster:

- [Configure LSF cluster components](#) on page 68
- [Apply a cluster configuration to an LSF cluster](#) on page 71

After selecting a cluster to configure, the following tabs represent the most commonly-used LSF configuration files that you can edit using Platform RTM:

- The Cluster Parameters tab corresponds with the `lsf.conf` file and allows you to edit the general LSF configuration to control the operation of LSF.
- The Cluster Hosts tab corresponds with the `lsf.cluster.cluster_name` file and allows you to edit the cluster configuration, the types of hosts in the cluster, and the resource assignments on the hosts in the cluster.
- The Resources tab corresponds with the `lsf.shared` file and allows you to define and edit common resource definitions that are shared by all defined clusters.
- The Queues tab corresponds with the `lsb.queues` file and allows you to define and edit the queues for this cluster.
- The Batch Hosts tab corresponds with the `lsb.hosts` file and allows you to configure the hosts, host groups, and job processing in this cluster. This tab has the following sub-tabs:
 - The Hosts sub-tab corresponds with the `Host` section of the `lsb.hosts` file and allows you to configure the hosts and job processing in the cluster.
 - The Host Groups sub-tab corresponds with the `Host Group` section of the `lsb.hosts` file and allows you to configure host groups in the cluster.
- The Batch Parameters tab corresponds with the `lsb.params` file and allows you to define general parameters to tune the operation of LSF, such as the timing within the system.
- The Batch Users tab corresponds with the `lsb.users` file and allows you to define users and user groups, along with the hierarchical fairshare and job slot limits for users and user groups. This tab has the following sub-tabs:
 - The Users sub-tab corresponds to the `User` section of the `lsb.users` file and allows you to define LSF users in the cluster.
 - The User Groups sub-tab corresponds to the `User Group` section of the `lsb.users` file and allows you to define LSF user groups in the cluster.
- The Resource Allocation tab corresponds with the `lsb.resources` file and allows you to define resource allocation and usage limits in the LSF cluster.

This tab displays the Limits sub-tab, which corresponds to the `Limit` section of the `lsb.resources` file and allows you to define limits for how many resources must be available for different classes of jobs to start, and which resource consumers the limits apply to.

When you save the configuration, Platform RTM adds the new parameters to the corresponding files in a section enclosed with "**#RTM Generated Section#**" text. Platform RTM backs up the original

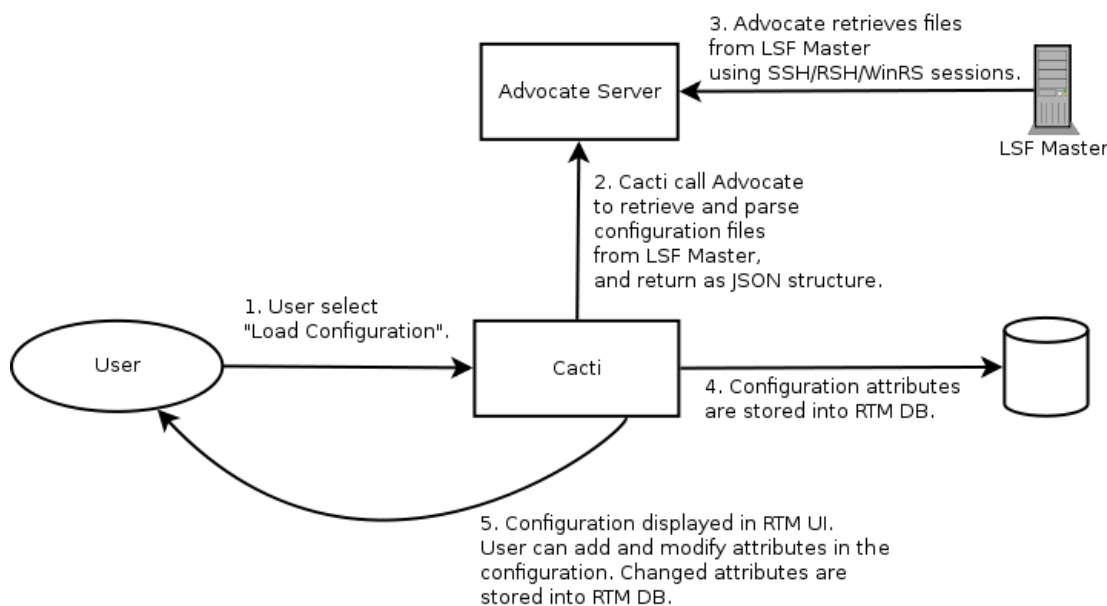
configuration by adding an "#RTM#" comment to the beginning of each line in the configuration file (which comments out that line in the file). Platform RTM then reconfigures the corresponding LSF cluster. If there are no errors with the cluster reconfiguration, Platform RTM backs up the new configuration file as a "known good" backup file.

Configure LSF cluster components

You can use the Platform RTM Console to create LSF cluster configurations that you can apply to the LSF cluster. Use the Platform RTM Console to load an LSF configuration, modify it, then apply the modified configuration to your cluster.

This allows you to create custom configurations that you can apply to the cluster at any time to change its configuration.

The following diagram describes the process under which Platform RTM loads the LSF cluster configuration:



1. Click the Config tab.
2. Under the Grid Management section of the Config menu bar, click LSF.

The Grid LSF Configuration page (with the LSF Configuration(s) tab) displays.

- The LSF Configuration(s) tab displays a list of previous LSF cluster configurations that you created using Platform RTM.
 - Clicking the Audit Logs tab displays a log of actions that users took in the Grid LSF Configuration page.
3. If you are not editing a previous LSF cluster configuration in the list, load the current LSF configuration from an existing cluster.

If you are editing a previous LSF cluster configuration in the list, skip this step.

- a) From the LSF Configuration(s) tab, click Load Configuration From Cluster.

The Load LSF Configuration(s) page displays.

- b) For the Cluster Name field, select the name of the cluster from which to load the current configuration and click Continue.
- c) At the confirmation prompt, click Yes to load the configuration from the specified cluster.

Platform RTM attempts to load the configuration files from the specified cluster. If successful, the name of the configuration defaults to "*cluster_name_default_indexnumber*".

This new configuration is accessible from the LSF Configuration(s) tab.

4. From the LSF Configuration(s) tab, click the name of the LSF configuration that you wish to edit.
5. For each component that you want to edit, click its corresponding tab.

- To edit the general LSF configuration to control the operation of LSF (the *lsf.conf* file), click the Cluster Parameters tab.

The Cluster Parameters tab displays a list of the parameters and corresponding values from the *lsf.conf* file.

- To edit the cluster configuration, the types of hosts in the cluster, or the resource assignments on the hosts in the cluster (the *lsf.cluster.cluster_name* file), click the Cluster Hosts tab.

The Cluster Hosts tab displays a list of hosts in the cluster.

- To edit the resource definitions for your cluster (the *lsf.shared* file), click the Resources tab.

The Resources tab displays a list of resources available to the cluster.

- To define or edit the batch queues in the cluster (the *lsb.queues* file), click the Queues tab.

The Queues tab displays a list of queues in the cluster.

- To edit the hosts and job processing configuration in the cluster (the *lsb.hosts* file), click the Batch Hosts tab, then click the Hosts sub-tab or Host Groups sub-tab.

- The Hosts sub-tab displays a list of hosts in the cluster.

- The Host Groups sub-tab displays a list of host groups in the cluster.

- To define general parameters to tune the operation of LSF, such as the timing within the system (the *lsb.params* file), click the Batch Parameters tab.

The Batch Parameters tab displays a list of the parameters and corresponding values from the hosts in the cluster.

- To edit the users, user groups, hierarchical fairshare, and job slot limits in the cluster (the *lsb.users* file), click the Batch Users tab, then click the Users sub-tab or User Groups sub-tab.

- The Users sub-tab displays a list of LSF users in the cluster.

- The User Groups sub-tab displays a list of LSF user groups in the cluster.

- To define limits for how many resources must be available for different classes of jobs to start, and which resource consumers the limits apply to (Limits section of the *lsb.resources* file), click the Resource Allocation tab.

The Resource Allocation tab displays the Limits sub-tab, which displays a list of resource allocation limits for the cluster.

6. If there are items that you need to copy, duplicate them from the list.

This option is not available in some tabs.

- a) Click the checkbox at the right side of each item that you want to delete.
- b) In the Choose an action field, select Duplicate and click Go.

7. If there are any items that you do not need, delete them from the list.

This option is not available in some tabs.

- a) Click the checkbox at the right side of each item that you want to delete.
 - b) In the Choose an action field, select Delete and click Go.
8. Add or edit an item for the LSF cluster component.

The method of adding or editing an item depends on the tab you are in, and some of the following options are not available in some tabs:

- To add a new item, click Add.
- To edit an existing item, click the name of the item that you want to edit.
- To edit the entire list of parameter values, click Edit.

Refer to *Platform LSF Configuration Reference* for more details on the parameters for the configuration file representing the component that you want to edit.

9. Specify (or update) the fields describing this item.

- Where available, the Advanced Attributes field allows you to specify other parameters from the corresponding configuration file.

Refer to *Platform LSF Configuration Reference* for more details on the parameters for the configuration file representing the component that you want to edit.

- For the Key field, specify the name of the parameter.
- For the Value field, specify the value that you want to assign to the parameter.
- Click Add Advanced Attributes to add more parameters to specify.

Note:

Platform RTM will not allow you to specify the following parameters as advanced attributes:

- LSF_LIM_PORT
- LSF_RES_PORT
- LSF_MBD_PORT
- LSF_SBD_PORT
- LSF_ENVDIR
- LSF_SHAREDIR

Platform RTM will check that you did not specify these parameters, but does not verify any other fields or advanced attributes before you save the configuration. Ensure that the values you specify for the fields and advanced attributes are valid for your cluster.

- Certain fields have the action icon to create a time-based configuration:



Click this action icon to display the fields to create a time-based configuration. This allows you to specify values based on different time windows. You can add additional time windows or delete pre-existing time windows. You can also disable time-based configuration for the specified field by clicking the time-based configuration action icon again.

- Time windows shown as tabs represent multiple parameter-value assignments that apply during the specified time window.

For example, in the Queues tab, you can specify values of the QJOB_LI MI T parameter (the total number of job slots that the specified queue can use) to change according to the time window.

You can also add new time window tabs or delete currently-existing time window tabs.

- Time windows shown as rows represent a single parameter-value assignment that applies during the specified time window.

For example, in the Hosts sub-tab of the Batch Hosts tab, you can specify multiple advanced attributes that apply to the specified host to change according to the time window.

You can also add new time window rows or delete currently-existing time window rows.


10. Click Save to save the settings for your item.

After configuring LSF cluster components, you can view audit logs of these changes. In the Grid LSF Configuration page, click the Audit Logs tab to view a log of actions that you took, with the outcomes of these actions.

Apply a cluster configuration to an LSF cluster

Apply a previous LSF cluster configuration that you created using the Platform RTM Console to the LSF cluster.

The following diagram describes the process under which Platform RTM applies the LSF cluster configuration:



1. Click the Grid tab.
2. Under the Grid Management section of the Grid menu bar, click LSF.
The Grid LSF Configuration page displays.
3. Click the checkbox at the right of the configuration that you want to apply.
4. In the Choose an action field, select Apply Configuration and click Go.
5. In the Apply Configuration dialog, specify the password for the LSF administrator of the cluster and click Continue.

The Grid LSF Configuration page displays the outcome of your action.

After applying the cluster configuration, you can view audit logs of these changes. In the Grid LSF Configuration page, click the Audit Logs tab to view a log of actions that you took, with the outcomes of these actions.

LSF log files

Platform RTM allows you to view the LSF log files in LSF hosts by using the Platform RTM Console.

Navigate to the LSF Logs page for a specific LSF host by clicking **By Host** under the Management section of the Grid menu bar, then clicking the **View Host Logs** action icon for the specific LSF host. This page allows you to view or download the LSF logs, and also allows you to configure the log levels for the specified host.

For more information on LSF logs, refer to *Error and Event Logging* in the *Administering Platform LSF* guide.

The following are actions you can take to work with the LSF log files:

- [View LSF log files](#) on page 72
- [Configure the LSF log level](#) on page 72

View LSF log files

1. Click the Grid tab.
2. Under the Management section of the Grid menu bar, click **By Host**.

The Host Control page displays a list of hosts in the clusters that Platform RTM monitors.

3. Navigate to the host for which you want to view the host logs and click the **View Host Logs** action icon, located in the Actions column.

The LSF Logs page displays a list of available LSF logs for the specified host.

4. Navigate to the LSF log file that you want to view and click the action icon corresponding with the log action that you want to perform:

Note:

Make sure that the LSF log and conf directories are shared. You can do **Head**, **Tail**, and **Download Host** actions only if those directories are shared.

If the LSF log and conf directories are not shared, then you cannot read the log files directly but you can download the files using the **Download Host** action.

- **Head Host Logs** displays the earliest entries for the specified log file.
- **Tail Host Logs** displays the latest entries for the specified log file.
- **Download Host Logs** allows you to download the entire log file onto your local file system.

Configure the LSF log level

Change the level of events that are recorded in the log files.

1. Click the Grid tab.
2. Under the Management section of the Grid menu bar, click **By Host**.

The Host Control page displays a list of hosts in the clusters that Platform RTM monitors.

3. Navigate to the host for which you want to view the host logs and click the **View Host Logs** action icon, located in the Actions column.

The LSF Logs page displays.

4. Click Configure Log Level.

The LSF Configure Log Level page displays.

5. Change the LSF log level for the LSF daemons.
 - To reset the log level for all the LSF daemons to the starting state, click Reset for All Daemons, then click OK.
 - To change the current log level, specify the LSF daemon, desired log level, and log classes, then click Apply.

Repeat this for each LSF daemon that you want to configure.

Note:

You can only configure the log level for the `mbat chd` and `mbschd` daemons if you selected the master host.

High availability applications

For more information on monitoring high availability applications, refer to [Monitor high availability applications](#) on page 43.

Navigate to the HA Configuration page by clicking HA under the Grid Management section of the Config menu bar. This page allows you to define and configure high availability applications that can run on clusters that Platform RTM monitors.

Important:

You must only use the Platform RTM Console to configure high availability applications. If you manually edit the high availability application configuration files, you may encounter configuration conflicts.

- [About high availability applications](#) on page 74
- [Manually start a high availability application on the primary or failover host](#) on page 76
- [Create or edit a high availability application configuration](#) on page 76
- [Apply a high availability application configuration to the LSF cluster](#) on page 78
- [Administer high availability application configurations](#) on page 78

About high availability applications

You can use Platform RTM to set up the LSF cluster to support and manage high availability (HA) applications.

High availability applications address the following requirements:

1. A system can monitor the health of applications and automatically restart these applications on available hosts if there are failures.

High availability applications restart on the failover host if the primary host unexpectedly goes down.

2. Client hosts can access applications without knowing the physical locations in which the applications are running.

Platform RTM and LSF do not directly address this requirement. A hardware load balancer or a software solution can provide a single, virtual IP address that automatically redirects to the current physical location of the high availability application.

Failover behavior

High availability applications are configured with a primary host and a failover host. The application normally runs on the primary host, but in the event of an outage, the application will instead run on the failover host.

Platform LSF first starts the application on the primary host. If the primary host goes down, LSF detects that the application in the primary host is non-responsive and restarts the application on the failover host.

Note that even if the primary host comes back up, the application running on the failover host will not fail back to the primary host unless the failover host goes down (in which case it will fail back to the primary host).

You can manually override this behavior by starting the application on the failover host if you are performing planned maintenance on the primary host, or by starting the application on the primary host if it is currently running on the failover host. For further details, refer to .

Note:

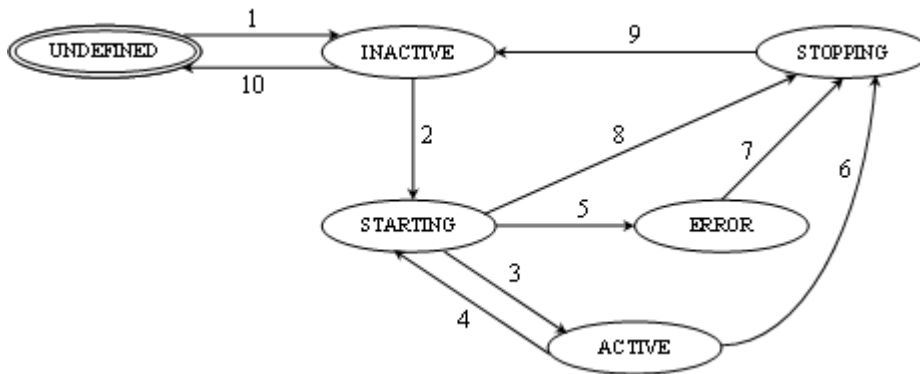
If the primary host is running and only if the application on the primary host unexpectedly goes down, Platform LSF tries to restart the application on the primary host.

High availability application state transition

High availability applications can have six different states: UNDEFI NED, I NACTI VE, STARTI NG, STOPPI NG, ACTI VE, and ERROR.

When a high availability application starts, it continues running unless there is a failure. That is, the start execution command does not run in the background.

The following diagram describes the state transition of high availability applications:



The following describes the status of the application during each state transition:

1. UNDEFI NED to I NACTI VE:

The application is created or updated in LSF using the Platform RTM HA Configuration page, if there are no errors in the configuration.

2. I NACTI VE to STARTI NG:

The application goes from the I NACTI VE state to the STARTI NG state under the following conditions:

- Immediately after the application is successfully created.
- Immediately after LSF restarts (for example, on a failover host).
- After you manually start up the application in the Applications dashboard in the Platform RTM Console.

3. STARTI NG to ACTI VE:

The application goes from the STARTI NG state to the ACTI VE state once LSF allocates all resources and starts all applications that this application depends on.

4. ACTI VE to STARTI NG:

The application goes from the ACTI VE state to the STARTI NG state if the application is terminated due to internal or external errors (for example, the application crashes or someone kills the application from outside of LSF).

By default, LSF attempts to restart the application (STARTI NG to ACTI VE) up to ten times.

5. STARTI NG to ERROR:

The application goes from the STARTI NG state to the ERROR state for the following reasons:

- LSF cannot find the applications that this application depends on.
 - LSF tried to restart the applications ten times (the default value).
6. ACTIVE to STOPPING:

The application goes to the STOPPING state if you manually stop the application from the Application dashboard in the Platform RTM Console.
 7. STARTING to STOPPING:
 8. ERROR to STOPPING:

The application goes to the STOPPING state if you manually stop the application from the Application dashboard in the Platform RTM Console.
 9. STOPPING to INACTIVE:

The application goes from the STOPPING state to the INACTIVE state after a predefined stop action is performed and LSF releases all resources associated with the application. Once the application is stopped, it will not start again until you manually start the application from the Application dashboard in the RTM Console or restart LSF.

By default, LSF attempts to restart the application (ACTIVE to STARTING) up to ten times.
 10. INACTIVE to UNDEFINED:

The application goes from the INACTIVE state to the UNDEFINED state if you removed the application from LSF using the Application dashboard in the Platform RTM Console.

Manually start a high availability application on the primary or failover host

Manually start a high availability application on the specified host (either the primary host or the failover host). This allows you to override the configured failover behavior of the specified application.

If you are planning to perform maintenance on the primary host that requires it to be down for a period of time, you can manually start the high availability application on the failover host by using the Migrate action.

If the high availability application is running on the failover host and you want it to run on the primary host again, you can manually start the application on the primary host by using the Restart action.

1. Click the HA tab.

The Applications Dashboard appears.
2. Select the checkboxes to the right of each high availability application that you want to manually start.
3. In the Choose an action field, select the action you want to take and click Go.
 - To manually start the application on the failover host, select Migrate.
 - To manually start the application on the primary host, select Restart.

Create or edit a high availability application configuration

You can use the Platform RTM Console to create high availability application configurations that you can apply to the LSF cluster. Use the Platform RTM Console to load a high availability application configuration, modify it, then apply the modified configuration to your cluster.

This allows you to create multiple groups of high availability applications configurations that you can apply to the cluster at any time.

1. If you are editing a high availability application configuration that is In Use, stop the applications that you want to edit.

- a) Click the HA tab.

The Application Dashboard displays.

- b) Click the checkbox at the right side of each high availability application that you want to edit.
- c) In the Choose an action field, select Stop and click Go.

The application may take a few minutes to stop. Verify that the application has stopped before proceeding to the next step.

2. Click the Config tab.

3. Under the Grid Management section of the Config menu bar, click HA.

The HA Configuration page displays a list of high availability application configurations available for your clusters.

4. If you are not editing a previous high availability application configuration in the list, load the current high availability application configuration either from an existing cluster or from an exported file.

If you are editing a previous high availability application configuration in the list, skip this step.

- To load the current high availability application configuration from an existing cluster, click Load HA Configuration From Cluster, then click Yes.
- To load the current high availability application configuration from an exported file, click Import HA Configuration From File, browse to the exported file, then click Import.

This new configuration is accessible from the HA Configuration page.

5. From the HA Configuration page, click the name of the high availability application configuration that you wish to edit.

The HA Applications section displays a list of high availability applications in the selected configuration.

6. If there are any high availability applications that you do not need, delete them from the list.

- a) Click the checkbox at the right side of each application that you want to delete.
- b) In the Choose an action field, select Delete and click Go.

7. Add or edit a high availability application.

- To add a new application, click Add.
- To add a new application by copying a previous application:
 1. Click the checkbox at the right side of an application to copy.
 2. In the Choose an action field, select Copy and click Go.
 3. At the copy confirmation prompt, click Yes.
 4. Specify a name for the new application and click Save.
 5. Click the name of the new high availability application to edit it.
- To edit an existing high availability application, click its name.

The Application Edit page displays.

8. Specify (or update) the fields describing this high availability application.

Important:

If you imported a configuration from an exported file, you need to make changes to the applications before applying the configuration to the cluster, including the primary host, failover host, the execution user name, and the execution user password.

9. Click Save or Create to save the settings for the high availability application.

If your high availability application configuration is not In Use (for example, if you imported a configuration from an exported file), you need to apply the configuration to the LSF cluster. For further details refer to [Apply a high availability application configuration to the LSF cluster](#) on page 78.

Apply a high availability application configuration to the LSF cluster

Apply a previous high availability application configuration that you created using the Platform RTM Console to the LSF cluster. This removes all existing high availability applications and replaces these with a new group of high availability applications.

1. Identify the applications that are modified or deleted and stop those applications that are running on LSF.

- a) Click the HA tab.

The Application Dashboard displays.

- b) Click the top-most checkbox at the column heading to select all applications.
- c) In the Choose an action field, select Stop and click Go.

The applications may take a few minutes to stop. Verify that the applications have stopped before proceeding to the next step.

2. Click the Config tab.
3. Under the Grid Management section of the Grid menu bar, click HA.

The HA Configuration page displays a list of high availability application configurations available for your clusters.

4. Click the checkbox at the right of the configuration that you want to apply.
5. In the Choose an action field, select Apply and click Go.
6. Follow the dialog prompts to apply the configuration to the cluster.

The HA Configuration page displays the outcome of your action.

Administer high availability application configurations

Use the Platform RTM Console to administer the high availability application configurations.

1. Click the Config tab.
2. Under the Grid Management section of the Grid menu bar, click HA.

The HA Configuration page displays a list of high availability application configurations available for your clusters.

3. Click the checkbox at the right of the configuration that you want to administer.
4. In the Choose an action field, select the action that you want to perform on the configuration and click Go.

- To rename the configuration, select Rename.
- To copy the configuration, select Copy.
- To export the configuration to an external file, select Export.
- To delete the configuration, select Delete.

Note:

Before applying a high availability configuration (by selecting **Apply**), you need to stop all modified high availability applications first. For further details refer to [Apply a high availability application configuration to the LSF cluster](#) on page 78.

5. Follow the dialog prompts to complete the action on the configuration.

Performance and maintenance

Database maintenance

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 83
- [Enable automatic data archiving](#) on page 84
- [Restore the Platform RTM system using the Platform RTM Console](#) on page 84

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 Config tab.
2. Under the Cacti Configuration section of the Config 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 Config 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 or .zip file contains the following files:
 - cacti_db_backup.sql
 - cacti_db_struct_backup.sql

- `rtm/etc/rtm.lic`
- `rtm/etc/lspollerd.conf`
- `rtm/etc/cluster_id/lsf.conf`
- `rtm/etc/cluster_id/ego.conf` (for LSF 7.x 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 Config tab.
2. Under the Cacti Configuration section of the Config 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 Config tab.
2. Under the Cacti Configuration section of the Config 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.

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 for SAS*.

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>

