

---

# Platform RTM Administrator Guide

Platform RTM  
Version 2.0.1  
Release date: February 2010



Copyright

© 1994-2010 Platform Computing Inc.

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

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

We'd like to hear from you

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

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

Document redistribution and translation

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

Internal redistribution

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

Trademarks

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

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

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

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

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

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

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

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

Third-party license agreements

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

---

# Contents

1	About Platform RTM .....	5
	Introduction to Cacti and Platform RTM .....	6
	User interface .....	7
2	Monitoring the cluster .....	13
	Viewing LSF cluster and job information .....	14
	Viewing threshold and alert information .....	26
	Viewing UNIX log file entries .....	27
	Viewing cluster, service, performance and host information .....	28
	Monitoring high availability applications .....	31
3	Administering Platform RTM .....	33
	Configuring cluster management .....	34
	Configuring cluster interaction .....	44
	Configuring date, time, and license information .....	52
	Working with thresholds and alerts .....	53
4	Administering Platform LSF using Platform RTM .....	57
	Cluster administration .....	58
	Controlling an LSF cluster .....	61
	Configuring an LSF cluster .....	63
	Working with high availability applications .....	67
5	Performance and maintenance .....	73
	Database maintenance .....	74
	Issues to consider .....	76



C H A P T E R

---

# 1

## About Platform RTM

# Introduction to Cacti and Platform RTM

## About this guide

Platform RTM caters to three user groups who are each responsible for HPC capacity management and planning: LSF Administrators, LSF users, and IT managers. This guide focuses on LSF administrators and LSF users.

Platform also provides RTM download, installation, and release information on *my.platform.com*.

For information specific to Cacti itself, see the Cacti documentation at <http://cacti.net/documentation.php>.

## About RTM and its interaction with Cacti

Platform RTM provides a rich graphical view of LSF clusters that has not been possible with any other product. RTM communicates the overall health of multiple LSF clusters, as well as data about past cluster performance.

RTM is based on Cacti, a widely popular Open Source product. Cacti was developed as an Open Source tool to provide IT administrators a way to graphically view the status of devices and services within their infrastructure. In recent years, with the release of the Cacti Plug-in Architecture, organizations using Cacti can now extend the Cacti framework to address other needs. Platform RTM is one such add-on to Cacti. RTM provides users the ability to view information about their LSF grids in a graphical way, and includes a near real-time reporting interface.

Out of the box, Cacti can monitor UPS devices, servers, services, databases, network switches, SANs and NASs. In addition, Cacti can record any time series data that can be obtained either through Net-SNMP or a script. Using this mechanism, Platform RTM provides the ability to view LSF data (such as execution hosts, users, queues, and job statistics), to control or configure LSF clusters (such as hosts, queues, and resources), and to monitor or configure high availability (HA) applications.

Together, Platform RTM and Cacti provide an opportunity for IT organizations to consolidate monitoring of an entire HPC computing infrastructure.

## Relationship between LSF and RTM

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

# 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

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, 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 RTM monitors.  This tab only exists on Linux versions of RTM.
Tab: HA	Opens the High Availability Application dashboard. View high availability (HA) applications and monitor application availability.
Tab: Settings	Allows you to customize either the layout of your graphs or the on-screen presentation of your grid.  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).  You may not see the Settings tab at all if the Cacti administrator has restricted your access.

Interface component	Description
Navigation bar	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.
Logout link	This link is to the right of the navigation bar. Click this link to log out of the system.

## 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 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>),

## Selection filtering

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



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.

## 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

## 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

## Page layout preferences and customizations

Depending on the configurations made by the 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 Setting 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.

The Settings page provides access to numerous other tabbed pages where you can configure, in detail, your display preferences.



## Monitoring the cluster

## Viewing 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 RTM and LSF, see [Configuring cluster interaction](#) on page 44.

For more information on configuring LSF clusters using RTM, see [Cluster administration](#) on page 58.

The Grid menu bar has seven sections: Dashboards, Management, Job Info, User/Group Info, Load Info, Host Info, and Reports. The following pages in these sections are specific to Platform RTM:

- **Dashboards: Cluster and Host pages**  
These pages provide an overview of LSF cluster and host health.
- **Management: By Cluster, By Host, By Queue, By Job**  
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).
- **Job Info: By Host, By Host Group, By Project, By Queue, By Array, Details pages**  
These pages provide information about LSF jobs, either at the level of host, host group, queue, or job group/job array. You can also view detailed information about specific jobs.
- **User/Group Info: User Viewer and User Groups Viewer pages**  
These pages provide information about LSF users and user groups.
- **Load Info: Host and Host Group pages**  
These pages provide information about host load and host groups.
- **Host Info: Servers, Clients, and Groups pages**  
These pages provide information about LSF cluster servers, clients and host groups.
- **Reports: Daily Statistics and Parameters pages**  
These pages provide information about daily statistics and job parameters. You can filter statistics and batch system parameters for specific information.

## 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).

## 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.

## Cluster Details page

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




This page shows the following status information represented by icons:

**Table 1: LIM Status**

Icon	Status
	<ul style="list-style-type: none"> <li>• ok</li> </ul>
	<ul style="list-style-type: none"> <li>• locked</li> <li>• locked-user</li> <li>• locked-window</li> <li>• locked-master</li> </ul>
	<ul style="list-style-type: none"> <li>• busy</li> </ul>

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

**Table 2: Batch Status**

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

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

## Job Info section

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

### By Host page

Navigate to the By Host page 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.
- **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.

## By Host Group page

Navigate to the By Host Group page 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.
- **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.

## By Project page

Navigate to the By Project page 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.

## By Queue page

Navigate to the By Queue page 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.

## By Array page

Navigate to the By Array 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.

## Details page

Navigate to the Details 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. In some cases, the availability of this option is dependent upon the specified cluster name.

The screenshot shows the 'Batch Job Filters' panel with the following settings: Cluster: cluster, User: All, UGroup: All, Status: ACTIVE, Effic: All, Queue: All, Host: All, HGroup: All, Records: 30, Except: N/A, Resource: [BUILTIN] N/A, JobID: (empty), Search: (empty), Dynamic: checked, Cluster TZ: unchecked, and buttons for go, clear, and export.

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.

## 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 page

Navigate to the Host Load page 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.
- **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 page

Navigate to the Group Load page 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 `lsl load`.
- 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).

## Host Info section

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

## Servers page

Navigate to the Servers page 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.

## Clients page

Navigate to the Clients page 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.

## Groups page

Navigate to the Groups page 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.

## Daily Statistics page

Navigate to the 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 Grid page)

### General tab

Click the General tab to open the Grid Settings (General) page and define the default page to display, as well as the default LSF cluster for which you want to filter data.

The fields on this page are as follows:

- Your Main Screen. The default information page that opens when you click the Grid tab.
- Default Cluster. The default cluster name used in the filter on all RTM pages.
- Default Cluster Timezone. Whether to display job event times using the time zone of the cluster, or the RTM server.
- Default Job Status. The default setting used to initially filter the list of displayed jobs (on the Grid page, within Job Info > Details).
- All Job Status. Allows the user to choose the "ALL" option to show a list of all jobs (on the Grid page, within Job Info > Details).
- Show Inactive Users. Preference to show all LSF users, even if they have not run any jobs recently (normally they are not shown).
- Audible Alerting. Preference for an audible alert to system administrators when an LSF host becomes unavailable or unreachable.

- **Blink When Down.** Preference for a host icon to blink when that host is down.
- **Support Advanced Popup.** Enables the display of additional popup content including additional job links. This does not work with Internet Explorer (version 7 and earlier).

## Visual tab

Click the Visual tab to open the Grid Settings (Visual) page and 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.

- **Icon Size Transition Host Count.** When this number is exceeded, RTM uses smaller icons to represent hosts on the Dashboards > Host page.
- **Host Status Popup Transition Host Count.** When this number of hosts is exceeded, RTM disables Javascript popups to allow the screen to refresh faster.
- **Number of Records to Display.** When viewing records in tabular form throughout the RTM interface, you can display varying amounts of data. This setting provides each interface a default number of records to display.

## Timespans tab

Click the Timespans tab to open the Grid Settings (Timespans) page and control how your grid summary icons appear. Each user can display thresholds that they believe represent how their grid hosts are behaving.

This page contains the following fields:

- **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”.
- **Default Graph View Timeshift.** The default time shift when viewing job details.
- **Allow Graph to extend to Future.**
- **First Day of the Week.** Used for the dayshift timespan.
- **Start of Daily Shift.** Used for the dayshift timespan.
- **End of Daily Shift.** Used for the dayshift timespan.

## Clusters tab

Click the Clusters tab to open the Grid Settings (Clusters) page and customize your view of the Dashboards > Cluster page. Choose to show or hide various information columns.

## Queues tab

Click the Queues tab to open the Grid Settings (Queues) page and show or hide certain fields on the Job Info > By Queue page.

## Jobs tab

Click the Jobs tab to open the Grid Settings (Jobs) page and show or hide certain fields on the Job Info > 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.

## Job Export tab

Click the Job Export tab to open the Grid Settings (Job Export) page and 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 tab

Click the Arrays tab to open the Grid Settings (Arrays) page and customize your view of the Job Info > By Array page. Choose those fields that are most relevant for your environment.

## Job Graphs tab

Click the Job Graphs tab to open the Grid Settings (Job Graphs) page and customize the display of your Job Info > Details graph according to your preference.

This page contains the following fields:

- **Graph Columns.** The number of columns on the page that RTM uses to display a job's graphs. Choose either 1 or 2 columns.
- **Width.** The width of a job's graphs, in pixels
- **Height.** The height of a job's graphs, in pixels.

You may also customize the colors that RTM uses when constructing job graphs. The drop-down lists are color-coded to facilitate your choice.

# Viewing threshold and alert information

Click the Alerts tab to view threshold and alert information.

For more information configuring thresholds, see [Working with thresholds and alerts](#) on page 53.

## 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.

# Viewing UNIX log file entries

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

The Syslogs page displays entries from the UNIX log files 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 RTM, edit the `/etc/syslog.conf` file on that host and add the IP address of the RTM host to the end of the file. For more information on configuring log files in the UNIX host, refer to the *README for Installing RTM for SAS on Linux*.

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.

# Viewing cluster, service, performance and host information

Click the Graphs tab to view a graphical representation of your cluster status and details in the Graphs page.

## 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 view



Click the Graphs tab, then the Preview view 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 Graph 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.

### General subsection

Configure general graph options and display formats.

You can configure the following items from this page:

- Default RRA
- Default View Mode
- Default Graph View Timespan
- Display Graph View Timespan Selector
- Default Graph View Timeshift
- Allow Graph to extend to Future
- First Day of the Week
- Start of the Daily Shift
- End of Daily Shift
- Graph Date Display Format
- Graph Date Separator
- Page Refresh

### Graph thumbnails subsection

Configure the size of the thumbnails used to represent your graphs.

You can configure the following items from this page:

- Thumbnail Height
- Thumbnail Width
- Thumbnail Column
- Thumbnail Selection

### Tree View Mode subsection

Configure the tree defaults and display when in tree view mode.

You can configure the following items from this page:

- Default Graph Tree
- Default Tree View Mode
- Dual Pane Tree Width
- Expand Hosts

### Preview Mode subsection

Configure the preview mode display.

You can configure the following items from this page:

- Graphs Per Page

## List View Mode subsection

Configure the list view mode display.

You can configure the following items from this page:

- Graphs Per Page

## Graph Fonts subsection

Choose whether to use your own custom fonts and font sizes or the system defaults.

You can configure the following items from this page:

- Use Custom Fonts

# Monitoring 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 RTM monitors.

For more information on high availability applications, refer to [Working with high availability applications](#) on page 67.

## 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 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.
  - **Start** - This will start the high availability application based on the application definition.
  - **Stop** - This will stop the high availability application.

Monitoring the cluster

## Administering Platform RTM

# Configuring cluster management

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

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 sections (and their corresponding pages) are specific to configuring how Platform RTM manages your clusters:

- Management: Graph Management, Graph Trees, Data Sources, Devices, Threshold pages  
These pages allow you to configure a number of settings related to monitoring your cluster using Platform RTM.
- Grid Management: Pollers, Clusters, LSF, HA, and Utilities pages  
These pages allow you to add LSF clusters and perform certain database administration functions.

## 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.

## Alerts Item page

Navigate to the Alerts 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
- Weekend exemption: Disable threshold alerts on weekends
- Disable restoration email: Disable threshold alerts when the threshold has returned to normal
- Reset acknowledgement: Reset acknowledgements when the threshold has returned to normal
- High/low threshold values
- Threshold type: High/low, baseline, or time based.
- Event triggering (Shell command): Specifies event trigger commands or shell scripts in the event of a breach.
  - High Threshold Trigger Command/Script: If the threshold is breached because the data source exceeds this value, the threshold triggers the specified command or shell script.
  - Low Threshold Trigger Command/Script: If the threshold is breached because the data source drops below this value, the threshold triggers the specified command or shell script.
  - Norm Threshold Trigger Command/Script: If the threshold is breached, then returns to normal, the threshold triggers the specified command or shell script.
- Event triggering (Grid administrator host level triggers): Specifies host-level actions in the event of a breach.
  - Host Level Action (High Threshold): If the threshold is breached because the data source exceeds this value, the threshold triggers the specified action on the host.
  - Host Level Action (Low Threshold): If the threshold is breached because the data source drops below this value, the threshold triggers the specified action on the host.
- Email message body: Email alert message content. This specifies the template that is used in alert email notifications for this threshold.

---

### Note:

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

- Syslog settings
- Data type: Special formatting for the given data
- Re-alert cycle: The amount of time the threshold repeats the alert, if it is still in breach.

- Notify accounts and extra alert emails: Email addresses to be notified when the threshold raises an alert

## Placeholder tags

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

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

Placeholder name	Tag for threshold name	Tag for alert email template	Description
Cluster ID	clusterid	<clusterid>	The ID of the cluster.
Cluster name	cluster_name	<cluster_name>	The name of the cluster.
Cluster LSF master	cluster_lsfmaster	<cluster_lsfmaster>	The name of the LSF master host for the cluster.
Cluster LSF version	cluster_version	<cluster_version>	The version of LSF running in the cluster.
Cluster LSF LIM port	cluster_limport	<cluster_limport>	The port number of LIM running in LSF on the master host.
Custom data value	custom_custom_field_name	<custom_custom_field_name>	The custom data value from the data source that is linked in this alert. For example, custom_percent, custom_status.
Host name	host_hostname	<host_hostname>	The host name of the device linked in this alert.
Host description	host_description	<host_description>	The host description of the device linked in this alert.
Threshold description	Not available	<DESCRIPTION>	The threshold description.
Threshold host name	Not available	<HOSTNAME>	The host name of the threshold.
Threshold trigger time	Not available	<TIME>	The time in which the threshold triggered this alert.
Threshold graph URL	Not available	<URL>	The link to the URL of the threshold graph.
Threshold current value	Not available	<CURRENTVALUE>	The current value of the data field being monitored by the threshold, at the time of the alert email.

Placeholder name	Tag for threshold name	Tag for alert email template	Description
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 RTM pollers. These pollers collect information from the LSF cluster; RTM uses this data to build various reports for 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 RTM poller (for example, `/opt/rtm/l sf v e r s i o n / b i n` for Linux or `C: \Program Files\Platform Computing\RTM\l sf v e r s i o n \ b i n` 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 RTM, see [Cluster administration](#) on page 58.

- **Add.** Add a cluster for 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.
- **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.)
- **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 RTM Console.

For more details on configuring an LSF cluster, refer to [Configuring an LSF cluster](#) on page 63.

## 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 RTM monitors.

For more details on configuring high availability applications, refer to [Working with high availability applications](#) on page 67.

## Utilities page

Navigate to the Utilities page by clicking Utilities under the Grid Management section of the Config menu bar. This page shows information about 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 RTM database tables. See the Appendix 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 RTM creates at least one backup file.

## Cacti Configuration section—Grid Settings page

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

Contact Platform Computing for assistance in determining the optimal settings for your clusters.

### General tab

Click the General tab to open the Grid Settings (General) page and configure the default user settings for your cluster.

You can configure the following on this page:

- **Domain names to strip from the display**  
This 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 hostnames to substitute**
- **Minimum user screen refresh interval**  
Restrict 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**  
Restrict 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 as the job zoom function is system-intensive. You should set a smaller window for larger clusters to reduce system load.

- User group filter operation

Specify how your cluster handles user group filtering:

- User Group Membership: User accounts are assigned to a user group.
- Job Specification: Jobs are assigned to user groups at job submission time by using `bsub -G`.
- Maximum export rows

Restrict the maximum number of rows that your users can export to increase system performance. You should set less rows to increase system performance.

- Enable cluster CPU factor leveling

---

**Important:**

Do not enable this setting unless you understand how to apply CPU factoring to hosts in your cluster.

- Add a summary device using the host template grid summary.

All graphs for this summary device are automatically added as well.

## Poller tab

Click the Poller tab to open the Grid Settings (Poller) page and configure poller defaults for data collection, interval settings, and thresholds.

You can configure the following on this page:

- Enable daemons
- LSF poller mode
- Enable host auto purge

If you enable this, server hosts are automatically purged from the RTM database when they are removed from clusters.

## Maint tab

Click the Maint tab to open the Grid Settings (Maint) page and configure system maintenance settings. You can keep more data for smaller clusters, because there are less records for these clusters.

You can configure the following on this page:

- Time when past database records are removed
- Retention period for job details

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.

- Retention period for individual job records

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.

- Retention period for daily summary statistics

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.

- Maximum number of database records to remove
- Maximum down-time for daemons disabled for maintenance purposes
- Enable database backup

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.

- Database backup schedule
- Number of database backups to maintain
- Database backup file location
- Restore the database from the archives
- Database optimization schedule

MySQL performs database optimization according to this schedule.

## Archiving tab

Click the Archiving tab to open the Grid Settings (Archiving) page and configure database archiving settings.

Data archiving allows deep dive analysis that will not impact the system, because you can perform this analysis on the archive database instead of on a database that is currently in use.

You can configure the following on this page:

- Enable data archiving
- Frequency of data archiving
- Database type that will store data archives
- Name of host receiving data archives
- Name of the database receiving data archives
- Database account user name, password, and port for connecting with the database
- Enable RRD file creation for archiving during record purging

---

**Note:**

Enabling this will result in very large data archives.

- Storage location of archived RRD files

## Paths tab

Click the Paths tab to open the Grid Settings (Paths) page and configure cluster directories and file paths.

You can configure the following on this page:

- Location of log files on poller hosts (for example, `/var/www/html/cacti/log/cacti.log`). If the directory is found and verified, the message [OK: FILE FOUND] appears below this field.
- Location of job usage RRD and image files (for example, `/opt/cacti/gridcache`). If the directory is found and verified, the message [OK: DIR FOUND] appears below this field.

## Alerts tab

Click the Alerts tab to open the Grid Settings (Alerts) page and configure cluster threshold settings, including thresholds to identify when resources are idle/closed, low, busy, or starved.

You can configure the following on this page:

- When to consider a host “idle” or “closed” (%), based on job slots
- When to consider a host “idle with jobs” (%), based on CPU percentage
- When to consider a host “low on resources” (%), based on load average
- When to consider a host’s low physical memory “urgent” (%)
- When to consider a host’s low swap memory “urgent” (%)
- When to consider a host’s low temp memory “urgent” (%)
- The point at which an IO rate is high enough to consider a host “low on resources” (Kb/sec)
- The point at which a paging rate is high enough to consider a host “low on resources” (pages/sec)
- When to consider a host “busy” (%), based on CPU
- When to consider a host “busy” (%), based on the load average
- The point at which an IO rate is high enough to consider a host “busy” (Kb/sec)
- The point at which a paging rate is high enough to consider a host “busy” (pages/sec)
- When to consider a host “idle” (%), based on comparisons of load vs. running jobs (shows if a host has orphaned or non-cpu intensive jobs running)
- When to consider a host “idle” (%), based on a comparison of load vs. running jobs (shows if a host may be running jobs outside the grid management system)
- When to consider a host starved

## Aggregation tab

Click the Aggregation tab to open the Grid Settings (Aggregation) page and configure default behavior for project information aggregation, host information aggregation, and memory tracking.

You can configure the following on this page:

- Wallclock calculation method  
Set this field for chargeback calculations, depending on whether you charge for suspend time.
- Enable the tracking of submitted project names  
Enabling this allows you to collect job data based on project names.
- Track where jobs are submitted from
- Indicate aggregation method for collected project names  
Project aggregation is used for project names that contain hierarchical metadata to assist with tracking.
- Starting string position (digit) during collection
- Delimiter for separating hierarchy levels
- Number of significant delimiter fields
- Track job memory resources

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.

---

## Status/Events tab

Click the Status/Events tab to open the Grid Settings (Status/Events) page and configure default behavior, thresholds, and visual cues for job flapping, cluster and job efficiency, PID levels, and job dependencies.

If you change any of the colors to None, the corresponding event is not shown in the legend. All of these events are logged to the Cacti log regardless of whether you enabled them in this tab.

You can configure the following on this page:

- Track and highlight jobs that change state frequently (job flapping)
- Set job flapping threshold
- Color to indicate warning state for job flapping
- Track and highlight jobs that violate the job efficiency threshold
- Set the job start window

Platform RTM only tracks job running for at least this period of time.

- Set the job warning threshold for entire cluster
- Color to indicate warning state for job efficiency
- Set the job alarm threshold for entire cluster
- Color to indicate alarm state for job efficiency
- Number of warning/alarm events before issuing a corresponding message
- Number of clear events prior to issuing a NOTICE event message
- Track PID levels and generate log message if threshold exceeded
- Set PID threshold
- Highlight pending jobs with dependencies
- Colors to indicate different types of jobs

# Configuring 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 section (and its corresponding pages) is specific to configuring Grid and Cacti settings:

- Cacti Configuration: Cacti Settings, Grid Settings pages

These pages allow you to configure a number of settings related how Platform RTM interacts with your cluster.

## Grid Settings page

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

Contact Platform Computing for assistance in determining the optimal settings for your clusters.

## General tab

Click the General tab to open the Grid Settings (General) page and configure the default user settings for your cluster.

You can configure the following on this page:

- Domain names to strip from the display

This 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 hostnames to substitute
- Minimum user screen refresh interval

Restrict 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

Restrict 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 as the job zoom function is system-intensive. You should set a smaller window for larger clusters to reduce system load.

- User group filter operation

Specify how your cluster handles user group filtering:

- User Group Membership: User accounts are assigned to a user group.
- Job Specification: Jobs are assigned to user groups at job submission time by using the LSF `bsub -G` command.
- Maximum export rows

Restrict the maximum number of rows that your users can export to increase system performance. You should set less rows to increase system performance.

- Enable cluster CPU factor leveling

---

**Important:**

Do not enable this setting unless you understand how to apply CPU factoring to hosts in your cluster.

---

## Poller tab

Click the Poller tab to open the Grid Settings (Poller) page and configure poller defaults for data collection, interval settings, and thresholds.

You can configure the following on this page:

- Enable daemons

## Maint tab

Click the Maint tab to open the Grid Settings (Maint) page and configure system maintenance settings. You can keep more data for smaller clusters, because there are less records for these clusters.

You can configure the following on this page:

- Time when past database records are removed
- Retention period for job details

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.

- Retention period for individual job records

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.

- Retention period for daily summary statistics

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.

- Maximum number of database records to remove
- Maximum down-time for daemons disabled for maintenance purposes
- Enable database backup

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.

---

- Database backup schedule
- Number of database backups to maintain
- Database backup file location
- Restore the database from the archives
- Database optimization schedule

MySQL performs database optimization according to this schedule.

## Archiving tab

Click the Archiving tab to open the Grid Settings (Archiving) page and configure database archiving settings.

Data archiving allows deep dive analysis that will not impact the system, because you can perform this analysis on the archive database instead of on a database that is currently in use.

You can configure the following on this page:

- Enable data archiving
- Frequency of data archiving
- Database type that will store data archives
- Name of host receiving data archives
- Name of the database receiving data archives
- Database account user name, password, and port for connecting with the database
- Enable RRD file creation for archiving during record purging

---

### Note:

Enabling this will result in very large data archives.

- Storage location of archived RRD files

## Paths tab

Click the Paths tab to open the Grid Settings (Paths) page and configure cluster directories and file paths.

You can configure the following on this page:

- Location of log files on poller hosts (for example, `/var/www/html/cacti/log/cacti.log`). If the directory is found and verified, the message [OK: FILE FOUND] appears below this field.
- Location of job usage RRD and image files (for example, `/opt/cacti/gridcache`). If the directory is found and verified, the message [OK: DIR FOUND] appears below this field.

## Alerts tab

Click the Alerts tab to open the Grid Settings (Alerts) page and configure cluster alert and threshold settings, including alerts and thresholds to identify when resources are idle/closed, low, busy, or starved.

You can configure the following on this page:

- When to consider a host “idle” or “closed” (%), based on job slots
- When to consider a host “idle with jobs” (%), based on CPU percentage
- When to consider a host “low on resources” (%), based on load average
- When to consider a host’s low physical memory “urgent” (%)
- When to consider a host’s low swap memory “urgent” (%)
- When to consider a host’s low temp memory “urgent” (%)
- The point at which an IO rate is high enough to consider a host “low on resources” (Kb/sec)
- The point at which a paging rate is high enough to consider a host “low on resources” (pages/sec)
- When to consider a host “busy” (%), based on CPU
- When to consider a host “busy” (%), based on the load average
- The point at which an IO rate is high enough to consider a host “busy” (Kb/sec)
- The point at which a paging rate is high enough to consider a host “busy” (pages/sec)

- When to consider a host “idle” (%), based on comparisons of load vs. running jobs (shows if a host has orphaned or non-cpu intensive jobs running)
- When to consider a host “idle” (%), based on a comparison of load vs. running jobs (shows if a host may be running jobs outside the grid management system)
- When to consider a host starved

## Aggregation tab

Click the Aggregation tab to open the Grid Settings (Aggregation) page and configure default behavior for project information aggregation, host information aggregation, and memory tracking.

You can configure the following on this page:

- Wallclock calculation method
  - Set this field for chargeback calculations, depending on whether you charge for suspend time.
- Enable the tracking of submitted project names
  - Enabling this allows you to collect job data based on project names.
- Track where jobs are submitted from
- Indicate aggregation method for collected project names
  - Project aggregation is used for project names that contain hierarchical metadata to assist with tracking.
- Starting string position (digit) during collection
- Delimiter for separating hierarchy levels
- Number of significant delimiter fields
- Track license project job performance
- Track job memory resources

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.

---

## Status/Events tab

Click the Status/Events tab to open the Grid Settings (Status/Events) page and configure default behavior, thresholds, and visual cues for job flapping, cluster and job efficiency, PID levels, and job dependencies.

If you change any of the colors to None, the corresponding event is not shown in the legend. All of these events are logged to the Cacti log regardless of whether you enabled them in this tab.

You can configure the following on this page:

- Track and highlight jobs that change state frequently (job flapping)
- Set job flapping threshold
- Color to indicate warning state for job flapping
- Track and highlight jobs that violate the job efficiency threshold
- Set the job start window

Platform RTM only tracks job running for at least this period of time.

- Set the job warning threshold for entire cluster
- Color to indicate warning state for job efficiency

- Set the job alarm threshold for entire cluster
- Color to indicate alarm state for job efficiency
- Number of warning/alarm events before issuing a corresponding message
- Number of clear events prior to issuing a NOTICE event message
- Track PID levels and generate log message if threshold exceeded
- Set PID threshold
- Highlight pending jobs with dependencies
- Colors to indicate different types of jobs

## 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 tab for each category of Cacti settings that you can change.

### General tab

Click the General tab to open the Cacti Settings (General) page and configure the default Cacti settings.

You can configure the following on this page:

- Event logging
- Poller-specific logging
- The version of the SNMP utility installed in the RTM host
- The version of the RRDTool utility installed in the RTM host
- SNMP default settings
- Whether RTM prompts the user before deleting items

### Paths tab

Click the Paths tab to open the Cacti Settings (Paths) page and configure Cacti directories and file paths.

You can configure the following on this page:

- Location of SNMP binary files on the RTM host. If the files are found and verified, the message [OK: FILE FOUND] appears below these fields.
- Location of the RRDTool binary file on the RTM host (for example, /user/bin/rrdtool). If the file is found and verified, the message [OK: FILE FOUND] appears below this field.
- Location of the RRDTool font file on the RTM host. If the file is found and verified, the message [OK: FILE FOUND] appears below this field.
- Location of the PHP binary file on the RTM host (for example, /usr/bin/php). If the file is found and verified, the message [OK: FILE FOUND] appears below this field.
- Location of log file on the RTM host (for example, /opt/cacti/log/cacti.log). If the file is found and verified, the message [OK: FILE FOUND] appears below this field.
- Location of the Spine binary file. If the file is found and verified, the message [OK: FILE FOUND] appears below this field.

### Poller tab

Click the Poller tab to open the Cacti Settings (Poller) page and configure poller defaults.

You can configure the following on this page:

- Enable polling

- Poller type
- Polling interval
- Cron or scheduled task interval
- Maximum concurrent poller processes
- Spine-specific execution parameters
- Method used to determine host availability (None, Ping, SNMP, or both)
- Ping settings
- Failure count: Number of polling intervals a host must be down before logging an error.
- Recovery count: Number of polling intervals a host must remain up before returning the host to an "up" status.

## Graph Export tab

Click the Graph Export tab to open the Cacti Settings (Graph Export) page and configure graph export settings.

You can configure the following on this page:

- Export method
- Presentation method
- Tree display/export settings
- Thumbnail settings
- Export directory path
- Local scratch directory path
- Export schedule
- FTP server options

## Visual tab

Click the Visual tab to open the Cacti Settings (Visual) page and configure Cacti display settings.

You can configure the following on this page:

- Graph display settings
- Maximum data query field length
- Graph creation settings
- Data source display settings
- Device display settings
- Log management settings
- RRDtool font settings
- Maximum number of rows to display on a single page for syslog events

This option only appears on Linux hosts.

## Authentication tab

Click the Authentication tab to open the Cacti Settings (Authentication) page and configure Cacti authentication settings.

You can configure the following on this page:

- Authentication method
- Name of the guest user for viewing graphs

- Name of the user that Cacti uses as a template for new users
- LDAP settings
- EGO authentication settings

## Boost tab

Click the Boost tab to open the Cacti Settings (Boost) page and configure Boost server settings.

You can configure the following on this page:

- On demand RRD update settings for Boost
- Boost server settings
- Image caching settings for Boost
- Path to the lock file for on demand RRD updating. This lock file assists the boost poller in detecting overrun conditions.
- Boost debug file path

## Alerting tab

Click the Alerting tab to open the Cacti Settings (Alerting) page and configure alert and cluster threshold settings.

You can configure the following on this page:

- Disable all thresholds
- Base URL of the Cacti server
- Maximum number of thresholds to display per page
- Enable logging of threshold failures
- Enable logging of threshold changes
- Default alerting options
- Default baseline options
- Emailing options

## Misc tab

This tab only appears on Linux hosts.

Click the Misc tab to open the Cacti Settings (Misc) page and configure syslog event settings.

You can configure the following on this page:

- Syslog page refresh interval
- Syslog event retention period
- Syslog event email settings

## RTM Plugins tab

Click the RTM Plugins tab to open the Cacti Settings (RTM Plugins) page and configure RTM plugin settings.

You can configure the following on this page:

- DNS suffix for the RTM server. This setting can be found in the `ssh` or `telnet` configuration of the RTM host and is only required if your web browser cannot resolve host names.
- `ssh` terminal window display settings

## Mail/DNS tab

Click the Mail/DNS tab to open the Cacti Settings (Mail/DNS) page and configure email and DNS settings for the Cacti server.

You can configure the following on this page:

- Emailing options
- 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.

---

- SMTP server options
- DNS options

# Configuring 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.

## Timezone tab

---

**Note:**

This tab does not exist on Windows versions of 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 RTM. If at least one of your LSF clusters uses a Platform licenses, you must obtain and use a Platform license for RTM.

# Working with 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 with thresholds and alerts:

- [Create a threshold to trigger alerts](#) on page 53
  - [Create a threshold from a graph template](#) on page 53
  - [Create a threshold from a host](#) on page 53
- [Modify threshold settings](#) on page 54
- [Delete thresholds](#) on page 54
- [Acknowledge alerts for a single threshold](#) on page 55
- [Acknowledge alerts for multiple thresholds](#) on page 55

## Create a threshold to trigger alerts

You can create thresholds using the following sources:

- [Create a threshold to trigger alerts](#) on page 53
- [Create a threshold from a host](#) on page 53

### 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 Clusters 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 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.

## Acknowledge alerts for a single threshold

Acknowledge the triggered alerts for a single threshold to prevent future email and syslog notifications.

1. Click the Alerts tab.

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

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

The acknowledge icon changes into the reset acknowledge icon.

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

## Acknowledge alerts for multiple thresholds

Acknowledge the triggered alerts for multiple thresholds to prevent future email and system log notifications.

1. Click the Config tab.

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

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

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

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



# Administering Platform LSF using Platform RTM

# Cluster administration

You can use RTM to add or configure LSF clusters that RTM monitors.

The following is a list of actions that you can take to add or configure LSF clusters:

- [Add or edit clusters to RTM](#) on page 58
  - [Add clusters to RTM using the RTM Console](#) on page 58
  - [Add clusters to RTM using a script](#) on page 58
- [Add the RTM host to the LSF cluster as an LSF client](#) on page 59

## Add or edit clusters to RTM

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

- [Add clusters to RTM using the RTM Console](#) on page 58
- [Add clusters to RTM using a script](#) on page 58

## Add clusters to RTM using the RTM Console

Use the RTM Console to add an LSF cluster to RTM.

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

The Clusters Console page displays.

3. 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.

The Cluster Edit page displays.

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 all fields in the User Authentication settings section.

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 created a new cluster, add the RTM host to the new LSF cluster as an LSF client.

## Add clusters to RTM using a script

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

1. In the command line, navigate to the `plugins/grid` subdirectory of the cacti installation directory.

For example,

Linux: **`cd /opt/cacti/plugins/grid`**

Windows (64-bit): `cd C:\Program Files (x86)\Platform Computing\RTM\wwwroot\Cacti\plugins\grid`

- 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

- 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 7.0.5 cluster named `maincluster` with `lsf.conf` located in `/share/lsf/conf`:

```
php -q grid_add_cluster.php --type=0 --pollerid=6 --cluster_name=maincluster --cluster_env=/
share/lsf/conf
```

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

- Add the RTM host to the new cluster as an LSF client.
- 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 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 RTM installer add the host to the LSF cluster, you need to manually add the RTM host to the cluster as an LSF client to give RTM access to LSF cluster data. You also need to do this to any LSF cluster that you did not use the RTM installer to add.

- Log into the LSF master host.
- If the LSF master host cannot resolve the 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 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.
- Edit the `lsf_cluster.cluster_name` file and add the RTM host to the Host section.
  - Reconfigure LIM and restart `mbatchd` to apply your changes to the cluster.

**lsadmin reconfig**

**badmin mbdrestart**

- 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 7869**

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

# Controlling an LSF cluster

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

Click the Grid tab and navigate to the Management section of the Grid menu bar to run grid control commands on an LSF cluster.

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

- Hosts

You can open or close hosts in LSF clusters.

- Queues

You can 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)
- `bswitch` (switches unfinished jobs from one queue to another)
- `bresume` (resumes suspended jobs)
- `brun` (forces jobs to run immediately)
- `bst op` (suspends unfinished jobs)
- `bkill` (sends signals to kill unfinished jobs) - on Linux hosts only
- `bkill -r` (forces a job kill)
- `bkill -s` (sends a specific signal to kill a job) - on Linux hosts only

## Run grid control commands on an LSF cluster

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

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

1. Click the Grid tab.
2. Under the Management section of the Grid menu bar, click the link corresponding with the type of grid control commands you want to run.
  - By Cluster: Control cluster-level components such as mbat chd, LIM, or RES.
  - By Host: Open or close hosts in the LSF cluster.
  - By Queue: Control queues in the LSF cluster.
  - By Job: Control jobs that are submitted to the LSF cluster.
3. Select the checkbox next to at least one item in the list for which you want to run the grid control command.
4. In the Choose an action field, select an action and click Go.
5. If the grid control requires additional information, specify this information in the displayed fields.
6. Click Yes to run the grid control command.

# Configuring an LSF cluster

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

- [Configure LSF cluster components](#) on page 63
- [Apply a cluster configuration to an LSF cluster](#) on page 65

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

- 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 and job processing in this cluster.

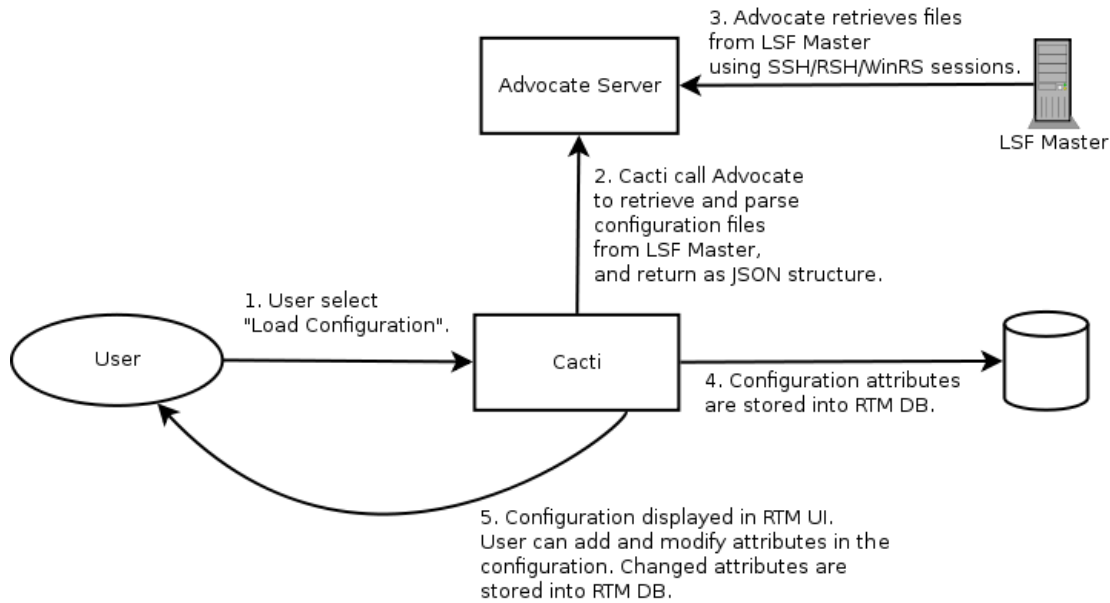
When you save the configuration, RTM adds the new parameters to the corresponding files in a section enclosed with "**#RTM Generated Section#**" text. RTM backs up the original configuration by adding an "**#RTM#**" comment to the beginning of each line in the configuration file (which comments out that line in the file). RTM then reconfigures the corresponding LSF cluster. If there are no errors with the cluster reconfiguration, RTM backs up the new configuration file as a "known good" backup file.

## Configure LSF cluster components

You can use the RTM Console to create LSF cluster configurations that you can apply to the LSF cluster. Use the 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 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 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 LSF Configuration Edit [new] page displays.

- b) For the Cluster ID field, select the name of the cluster from which to load the current configuration.
- c) Click Continue to load the configuration from the specified cluster.
- d) Specify the LSF administrator password for the cluster and click Continue.

RTM attempts to load the configuration files from the specified cluster. If successful, the name of the configuration defaults to "*cluster\_name\_default*".

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 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 Hosts tab.

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

6. If there are any items that you do not need, delete them from the list.
  - 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.
7. Add or edit an item for the LSF cluster component.

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

8. 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:**

RTM does not verify any 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.

---

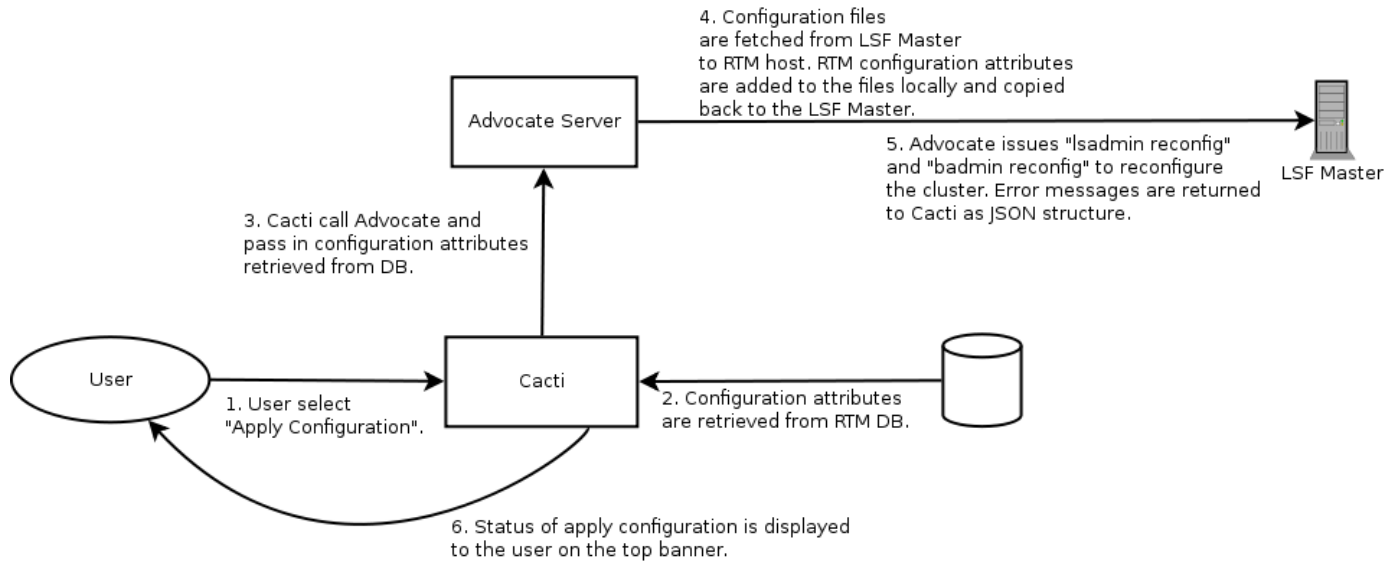
9. Click Save to save the settings for your item.

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 RTM Console to the LSF cluster.

The following diagram describes the process under which 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.

# Working with high availability applications

For more information on monitoring high availability applications, refer to [Monitoring high availability applications](#) on page 31.

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 RTM monitors.

---

## Important:

You must only use the 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 67
- [Manually start a high availability application on the primary or failover host](#) on page 69
- [Add or edit a high availability application](#) on page 69
- [Copy a high availability application](#) on page 70
- [Delete a high availability application](#) on page 70
- [Save high availability application configurations as a template](#) on page 71
- [View or apply existing high availability application templates](#) on page 71
- [Delete existing high availability application templates](#) on page 72

## About high availability applications

You can use RTM to set up the LSF cluster to support and manage high availability 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 application on the primary host unexpectedly goes down.

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

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 application on the primary host unexpectedly goes down (for example, 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).

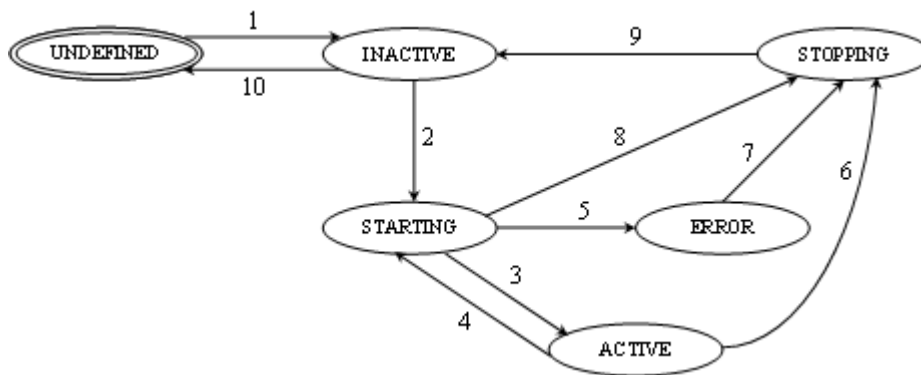
You can manually override this behavior by starting the application on the failover host if you are performing planned maintenance on the primary host, or by starting the application on the primary host if it is currently running on the failover host. For further details, refer to .

## High availability application state transition

High availability applications can have six different states: *UNDEFINED*, *INACTIVE*, *STARTING*, *STOPPING*, *ACTIVE*, 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. *UNDEFINED* to *INACTIVE*:

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

2. *INACTIVE* to *STARTING*:

The application goes from the *INACTIVE* state to the *STARTING* 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 RTM Console.

3. *STARTING* to *ACTIVE*:

The application goes from the *STARTING* state to the *ACTIVE* state once LSF allocates all resources and starts all applications that this application depends on.

4. *ACTIVE* to *STARTING*:

The application goes from the *ACTIVE* state to the *STARTING* 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 (*STARTING* to *ACTIVE*) up to ten times.

5. *STARTING* to *ERROR*:

The application goes from the *STARTING* 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*:
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 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 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.

## Add or edit a high availability application

Add or edit a high availability (HA) application to a cluster that RTM monitors.

1. If you are editing a high availability application, stop the application 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.
4. Add or edit a high availability application.
  - To add a new high availability application, click Add.
  - To edit an existing high availability application, click the name of the application that you want to edit.
5. Specify (or update) the fields describing your high availability application.
  - For the Execution Command [Start] field, specify the command to start the application. RTM does not verify the command that you specify here.  
This command must be able to run on both the primary and failover hosts, and cannot be a process that runs in the background. It needs to be a process that keeps running unless there is a failure that causes it to exit.
  - For the Execution Command [Stop] field, specify the command to stop the application. This command must be able to run on both the primary and failover hosts. If this field is unspecified, EGO will send a signal to kill the service in order to stop the application. RTM does not verify the command that you specify here.
  - For the Execution User and Execution User Password fields, specify the credentials for the user that will run the application. This user must exist on both the primary and failover hosts. RTM does not verify the user credentials that you specify here.
6. Click Create (or Save) to save the settings for your high availability application.
7. If you edited a high availability application, restart the application that you edited.
  - a) Click the HA tab.  
The Application Dashboard displays.
  - b) Click the checkbox at the right side of each high availability application that you edited.
  - c) In the Choose an action field, select Start and click Go.  
The application may take a few minutes to start.

## Copy a high availability application

Copy a high availability (HA) application.

This allows you to use the same configuration for a high availability application, and is useful when creating a new application with a similar configuration to one that you already use.

1. Click the Config tab.
2. Under the Grid Management section of the Config menu bar, click HA.  
The HA Configuration page displays.
3. Click the Copy application icon in the Actions column.  
The Application [copy - *application\_name*] page displays.
4. Specify the name of the new application and click Save to copy the application.

You can edit the new application to make further changes.

## Delete a high availability application

Delete a high availability (HA) application that you no longer need.

1. Stop the high availability application that you want to delete.
  - 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 delete.
  - 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.
4. Click the checkbox at the right side of each high availability application that you want to delete.
5. In the Choose an action field, select Delete and click Go.  
RTM displays a warning if you try to delete a high availability application that is currently running.

## Save high availability application configurations as a template

Save all high availability (HA) application configurations in a cluster as a template.

The high availability application configuration that you created for a cluster may be useful for other users, or as a way to organize your own high availability applications.

You can save the current configuration of all high availability applications in a cluster as a template to make available to other users, or to organize different high availability applications.

1. Click the Config tab.
2. Under the Grid Management section of the Config menu bar, click HA.  
The HA Configuration page displays.
3. Add, edit, or delete the high availability application configurations until you have all the applications configured correctly.
4. Click Save as template.  
The Template [new - *cluster\_name*] page displays.
5. Specify a name and description for your template and click Create to save your new template.

## View or apply existing high availability application templates

View a list of existing high availability (HA) application templates and optionally replace all existing applications in your cluster with the applications in a template.

To apply a template to a cluster, all high availability applications in that cluster must first be stopped.

When applying a template to a cluster, RTM removes all existing high availability applications from that cluster before replacing them with the applications in the template.

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

The HA Configuration page displays.

3. Click View templates.

The Saved Templates page displays.

4. To view the applications in a template, click the name of that template.

The Template [*template\_name* - *cluster\_name*] page displays.

5. If you want to apply the template to your cluster, click Apply.

RTM stops all existing high availability applications in that cluster before replacing them with the applications in the template.

## Delete existing high availability application templates

Delete existing high availability (HA) application templates that you no longer need.

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

The HA Configuration page displays.

3. Click View templates.

The Saved Templates page displays.

4. Click the checkbox at the right side of each template that you want to delete.
5. In the Choose an action field, select Delete and click Go.

## Performance and maintenance

# Database maintenance

Click the Config tab, navigate to the Cacti Configuration section of the Config menu bar, and click Grid Settings to perform or configure database maintenance using the RTM Console.

- [Backing up the database](#) on page 74
  - [Back up the existing database using the RTM Console](#) on page 74
  - [Enable automatic data archiving](#) on page 75
  - [Restore a backed-up database using the RTM Console](#) on page 75
  -

## Backing up the database

RTM allows you to back up and restore your configuration within the RTM Console.

The following files are backed up:

- `rtm.lic`: RTM license file. This file is not restored automatically when you restore your configuration within the 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 clusters only): The `ego.conf` file associated with each cluster.

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

- `grid_jobs`
- `grid_jobs_rusage`
- `grid_job_interval_stats`
- `poller_output`
- `poller_output_boost`

## Back up the existing database using the RTM Console

You must be using RTM versions 1.04 or later.

Prior to upgrading to a newer version of RTM or to a fully licensed version, you must first backup the existing database. Backing up your database is also recommended during scheduled server maintenance.

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:
  - Backup Cacti Database—Check this box to ensure the Cacti database is backed up when the maintenance script runs.
  - 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.)
5. Under the Grid Management section of the Config menu bar, click Utilities.
6. In the Database Administration section of the Utilities page, and 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 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 a backed-up database using the RTM Console

You must be using RTM versions 1.04 or later.

After you have completed any scheduled server maintenance, or finished upgrading your RTM version or license, you must restore the Cacti database that you previously backed up.

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

RTM requires access to the LIM (load information manager) port. If you do not specify this port, 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.

### Known issues

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

### Troubleshooting and FAQs

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