

Foglight for Amazon Redshift

Cartridge Guide

© 2023 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Table of Contents

Table of Contents	3
Introduction	5
Description	5
Business Challenge	5
Key Features	5
Foglight for Redshift Requirements	7
Installing and Configuring Agents	8
Redshift Pre-Configuration	9
IAM Policy Setup	9
IAM User Setup	9
Cartridge Installation	10
Creating and Configuring Agents	11
Using the Agent Installer Wizard (not yet available)	12
Using the Agent Status Dashboard	13
Agent Properties	14
Cluster Connection	14
Collection Intervals	14
DB Overrides	14
Statement Tracking	15
Options	15
Roles	16
Upgrading the Agent	17
Removing Monitored Databases (not yet available)	18
Administration	19
Opening the Databases Administration Dashboard (not yet available)	19
Reviewing the Administration Settings	19
Customizing Alarms for Foglight for Redshift Rules (not yet available)	20
Introducing the Alarms View	20
Modifying Alarm Settings	21
Reviewing Rule Definitions	24

Cloning Agent Settings	25
Configuring Email Notifications	26
Dashboards	29
Databases.....	29
Query Insights	29
Redshift Clusters	29
Cluster Overview.....	30
Nodes	31
Tables	32
Connections	32
WLM	33
Transactions	34
Queries.....	34
Configuration	35
Rules.....	36
Redshift Cluster Availability	36
Redshift CPU Utilization	36
Redshift Node Disk Space	36
Redshift Table Out of Date Statistics	36
Redshift Table Skew Value	36
Redshift Table Unsorted Percent	36
Redshift WLM Query Slots	36
Redshift WLM Queue Bottleneck.....	36
Reports.....	37
Redshift Cluster Summary.....	37
Redshift S3 Queries	37
Redshift Tables.....	37
Redshift Top Queries.....	37

Introduction

Description

Amazon Redshift is a fully-managed, highly-scalable database service on the AWS cloud. It can host data in a Redshift cluster and provide an interface to a data lake. Redshift can scale from a few hundred gigabytes of data to more than a petabyte.

Redshift uses columnar storage, data compression, and zone maps to deliver fast parallel processing. It also integrates machine learning and results caching to manage large workloads and improve response time with repeated workloads.

Business Challenge

While basic tools exist to monitor Redshift, these tools do not show relationships between data in a meaningful way, nor can they provide historical context. Amazon CloudWatch provides host performance metrics while Redshift operational data is available in system views, however neither interface provides a complete picture of your Redshift cluster. For environments with large clusters, the picture becomes even less clear.

Foglight for Redshift delivers a comprehensive overview of your Redshift clusters, combining CloudWatch metrics and query/load performance data with Host metrics. Combining this data provides a complete picture of your Redshift environment in a single, unified interface. No other tool provides this level of availability and performance monitoring with historical analysis. Foglight for Redshift allows DBA's to manage their environment in the same way they manage other legacy platforms. Foglight for Redshift provides the detailed information necessary to optimize query, table, and node performance.

Key Features

Foglight for Redshift brings true enterprise-class monitoring for Redshift to the industry's leading cross-platform database management solution. Leveraging both Redshift system data and the CloudWatch API, Foglight for Redshift combines performance and monitoring data into an industry-leading monitoring platform.

View all monitored Redshift clusters with health status, connections, workload, and critical performance metrics consolidated in a single dashboard. Drill down to a summary of cluster data with detailed information on nodes and tables including the key performance metrics and workload data necessary to ensure optimum performance.

Foglight for Redshift also provides information on WLM Queue states, active transactions, and performance of top queries. Users can quickly identify queries with long execution times or those that return a large amount of data or use excessive system resources. Queries to an S3 data lake made through Redshift are included in the performance monitoring.

Foglight for Redshift is preconfigured with features including out-of-the-box rules to provide automated notification of availability and performance and reports to provide regular updates to engineering and management. Foglight proactively creates performance baselines and notifies users of abnormalities.

Rules can be configured to trigger on pre-defined thresholds or on behavior that is unusual for a given timeframe. Notifications and reports are easily customized and tuned for your environment and specific business needs.

Foglight for Redshift Requirements

Foglight for Redshift is compatible with **Redshift 1.0.9865+**.

Foglight for Redshift can be installed on **FMS 5.9.2+** and agents require **FglAM 5.8.5.2+**.

Installing and Configuring Agents

Installation of Foglight for Redshift is covered in the following sections and should be performed in order:

- [Redshift Pre-Configuration](#)
- [Cartridge Installation](#)
- [Creating and Configuring Agents](#)

Redshift Pre-Configuration

In order to allow full monitoring of Redshift, the agent will require an IAM user with sufficient privileges to execute system queries as well as access to metrics through CloudWatch. Skip any steps in user or policy creation that have already been performed.

IAM Policy Setup

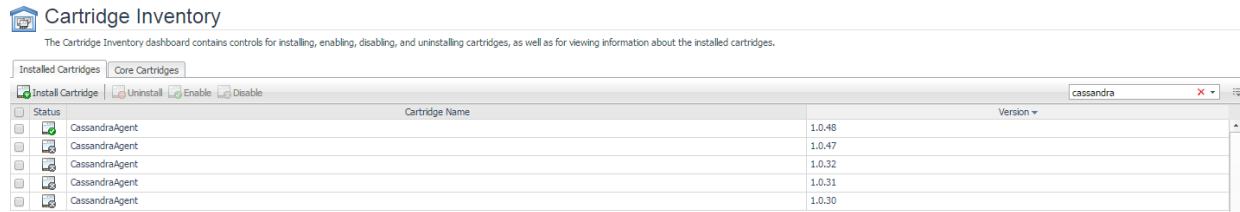
1. Using the IAM service, create a policy called *ClusterCredentialsPermission* and use the following configuration
 - a. Service – Select Redshift
 - b. Actions – In Access Level > Write select *GetClusterCredentials*. In Permissions Management select *CreateClusterUser* and *JoinGroup*
 - c. Resources – Select *All Resources*
2. Review policy and save

IAM User Setup

1. Using the IAM service, create an IAM user
2. Create a name for the user and select the “Programmatic Access Option” (uses access ID and secret access key)
3. When configuring permissions, select the “Attach existing policies directly” option and search for and add *CloudWatchReadOnlyAccess* and *ClusterCredentialsPermission*
4. Finish creating the user

Cartridge Installation

1. Open Foglight Management Console.
2. From the navigation pane, select: **Dashboards > Administration > Cartridges > Cartridge Inventory**. The Cartridge Inventory screen appears. For more information on agents, see the *Foglight User Guide*.
3. Load the **RedshiftAgent-xxxx.car** file by browsing to the location where the .car file exists and then clicking on “Install Cartridge”. Leave the “Enable on Install” check box checked.
4. Once the installation is completed on the Foglight Management Server, the Redshift Cartridge will appear in this list below as an installed cartridge.



Creating and Configuring Agents

Agents can be created in one of two ways:

- [Using the Agent Installer Wizard](#) (feature not yet available)
- [Using the Agent Status Dashboard](#)

The Agent Installer Wizard simplifies the agent creation and configuration process and can be accessed from the Databases dashboard. For advanced configuration or modification of agent properties post-creation, use the Agent Status dashboard.

Using the Agent Installer Wizard (not yet available)

Foglight for Redshift provides a graphic, intuitive method for creating and configuring agents, which can be used instead of Foglight's default method for creating agents and editing their properties using the Agent Status dashboard.

To run the instance installation wizard:

1. On the navigation panel, click Homes > Databases.
2. Click the Redshift box in the Databases View, and then click Monitor.
3. The Agent Installer Wizard dialog box appears.
4. The first card - Agent Deployment – has two fields:
 - a. Agent Name – Provide a name for the agent that will be created. This is not canonical and should be representative of the database instance that this agent will monitor.
 - b. Agent Manager - Choose the agent manager on which the agent should run. Considerations for this may include physical or virtual locality to the monitored instance, allocated resources, or grouping with other agents of the same type or monitored environment. If the agent package has not been deployed to this Agent Manager yet, it will be installed when the first agent of this type is created.
5. The second card – Agent Properties – requires a basic set of parameters for connecting to and monitoring the database instance. A full explanation of these properties is available in the [Agent Properties](#) section of this document.
6. The third card – Agent Summary – displays a review of the configuration that will be created and an option allowing the agent to be activated after creation. If the configuration looks good, click the Finish button to start the process.
7. When the process completes, a results screen will appear showing the results of agent creation. If the agent was not created, follow the instructions on the results screen. If successful, the database instance should appear in the Databases table within a few minutes.

Note: If the agent was created successfully but data is not appearing, go to the Dashboards > Administration > Agents > Agent Status page and click the icon in the Log File column for the agent you created. In most cases, the reason for the failure will be obvious. You can also refer to the *Foglight for Redshift Installation and Troubleshooting* document for common errors and solutions. If the solution requires reconfiguring the agent properties, follow steps 3-7 of the [Using the Agent Status Dashboard](#) section.

Using the Agent Status Dashboard

The Agent Status page can be used to create new agents and configure and manage existing agents. To access the page from the navigation pane, select: Dashboards > Administration > Agents > Agent Status.

Use the following steps to create a new agent instance:

1. If the Redshift agent package has never been deployed to the FglAM that will be used to host the agent, this must be done before an agent has been created. You can use the Deploy Agent Package button on the Agent Status or Agent Managers page to perform this.
2. Click the Create Agent button and follow the instructions for the cards:
 - a. **Host Selector** - Choose the Agent Manager on which the agent should run. Considerations for this may include physical or virtual locality to the monitored instance, allocated resources, or grouping with other agents of the same type or monitored environment.
 - b. **Agent Type and Instance Name** – Select the RedshiftAgent type. Then, select the Specify Name radio button and provide a name for the agent that will be created. This is not canonical and should be representative of the database instance that this agent will monitor.
 - c. **Summary** – Click Finish.
3. Once the agent has been created, click the checkbox next to the Redshift agent.
4. Click the **Edit Properties** button.
5. Select **Modify the default properties for this agent**.
6. Edit the agent properties for the Redshift agent instance:
 - [Cluster Connection](#)
 - [Collection Intervals](#)
 - [DB Overrides](#)
 - [Statement Tracking](#)
 - [Options](#)
7. Click the **Activate** button.

To modify the properties for an existing agent, skip to step 3 and Deactivate, then Reactivate the agents after changing the configuration.

Agent Properties

This is a full list and explanation of the configurable properties of the Foglight for Redshift agent. The Agent Installer Wizard provides access to the essential subset of available properties. To modify other properties or modify the agent configuration after creation, use the Agent Status dashboard.

Cluster Connection

The agent requires a connection to the cluster in order to gather information about the cluster and data structure. It is very important that the values provided here exactly match the details specified in the Cluster Properties of the AWS console for this Redshift cluster.

- **Cluster Name** – The name of the Redshift cluster. This is a unique key that identifies the cluster.
- **Region** – The EC2 Availability Zone that the cluster was created in. Described as Zone in the Cluster Properties.
- **IP or Hostname** – The Endpoint address for the cluster minus the port number.
- **Port** – The port number included in the Endpoint address.
- **Database** – Name of a database to connect to. All databases will be monitored, but data gathered from system tables and views available to the entire cluster will be through this database.
- **Username** – Redshift database user that must have at a minimum SELECT privilege on all system tables and views.
- **Access Key ID** – Access Key ID of the IAM user created during pre-configuration or existing user with necessary attached policies.
- **Secret Access Key** – Secret access key for the user described above.
- **SSL Mode** – Enforced security for making a connection to the Redshift cluster. If `require_ssl` is set to true for the cluster parameter group, this must be set to at least Enable.

Collection Intervals

The Collection Interval fields in the agent properties are used to set the sample frequencies. You can turn off a collection by setting the interval to 0. The defaults are set based on the type of data being collected for relevancy. Data gathered from CloudWatch will always be collected at 5-minute intervals.

- **DB List Refresh** – Checks existing databases in order to remove collection tasks for dropped databases or schedule new collections for newly created databases.
- **Availability** – Checks ability of the agent to connect to and request data from the cluster.
- **Cluster** – Retrieves cluster-level data from the database.
- **Tables** – Retrieves table structure and statistics from each database.
- **Configuration** – Retrieves configuration values from `pg_settings`.

DB Overrides

The DB override properties allow you finer-grained control over database collections. The default list is populated with system databases that have restricted use and will cause the agent to generate errors if access attempts are made. If a database is not specified here, it will use the parameters set in the Collection Intervals section.

- **Database Name** – Name of the database where collections are being overridden.
- **Ignore Tables** – If set to true, table information for this database will not be collected.
- **Interval** – This collection interval will override the default Tables setting in the Collection Intervals section for this database.

Statement Tracking

- **Enable Statement Tracking** – Enables or disables use of statement tracking.
- **# of Top Statements** – The maximum number of statements for the agent to collect during each sample period.
- **Order By** – The statements are sorted in order to gather the # of top statements specified in the previous property. If all statements are being gathered, this field is unimportant. Otherwise, several self-explanatory options are available based on existing or calculated column data in the view.
- **Enable S3 Query Tracking** – Enables statement tracking specifically for S3 queries.
- **# of Top S3 Queries** – The maximum number of S3 queries for the agent to collect during each sample period.

Options

- **FMS Time Offset** – This option allows conversion of timestamp information from various collections. Timestamp information from Redshift is generally in UTC time. The default value of SYSTEM will automatically convert timestamps based on the offset of the machine that hosts the agent from UTC.

Roles

Two roles, Redshift User and Redshift Administrator, are installed with the cartridge. Viewing Redshift dashboards requires that a user be assigned one of these or have the core Administrator role. The Redshift Administrator role does not currently grant any additional privileges.

Upgrading the Agent

1. Go to Dashboards > Administration > Cartridges > Cartridge Inventory and click the Install Cartridge button.
2. Locate the .car file on your system and install it with auto-enable selected. If you get a message that a bundled cartridge is of an older version than the one currently enabled on your FMS and will not be enabled, ignore it and continue.
3. Once the cartridge is installed and enabled, go to Dashboards > Administration > Agents > Agent Managers. Agent Managers that can be upgraded with newer agent packages will show “yes” in the Upgradable | Agents column. Select all Agent Managers you wish to upgrade and click the Upgrade button.

Note: If an Agent Manager is not upgradable, check that the Agent Manager version is compatible with the newer agent version. If it is not, the Agent Manager will need to be upgraded first.

Removing Monitored Databases (not yet available)

1. Go to the Databases dashboard.
2. Select the databases you wish to remove.
3. Click the Settings button, then click ok.

Note: Doing this will remove the monitoring agents as well as the historical data already collected. If you wish to delete only the agents, you can do that on the Administration > Agents > Agent Status page. Because the Databases dashboard only shows databases which are being actively monitored, you will only be able to view these databases by going directly to the Redshift dashboard.

Administration

Opening the Databases Administration Dashboard (not yet available)

You can edit agent settings for one or more Redshift instances on the Databases > Administration dashboard.

NOTE: If you attempt to select instances of more than one type of database, such as a Redshift database and an Oracle database, an error message is displayed.

To open the Databases Administration dashboard:

1. In the navigation panel, under **Homes**, click **Databases > Redshift**.
2. Select the check boxes beside one or more Redshift instances.
3. Click **Settings** and then click **Administration**. The Administration dashboard opens, containing settings for all the selected agents. Settings are broken down into categories, which are organized under a Redshift tree.

Reviewing the Administration Settings

The Databases Administration dashboard allows settings options for collecting, storing, and displaying data, which apply to all the currently selected agents. Click a category of settings on the left (for example: Connection Details) to open a view containing related settings on the right.

To view the full list of selected agents, click the **Selected Agents** button at the upper right corner of the screen. To change the list of agents to which the metrics will apply, exit the Databases Administration dashboard, select the requested agents and re-open the view.

Customizing Alarms for Foglight for Redshift Rules (not yet available)

Many Foglight for Redshift multiple-severity rules trigger alarms. To improve your monitoring experience, you can customize when alarms are triggered and whether they are reported. You can also set up email notifications.

Introducing the Alarms View

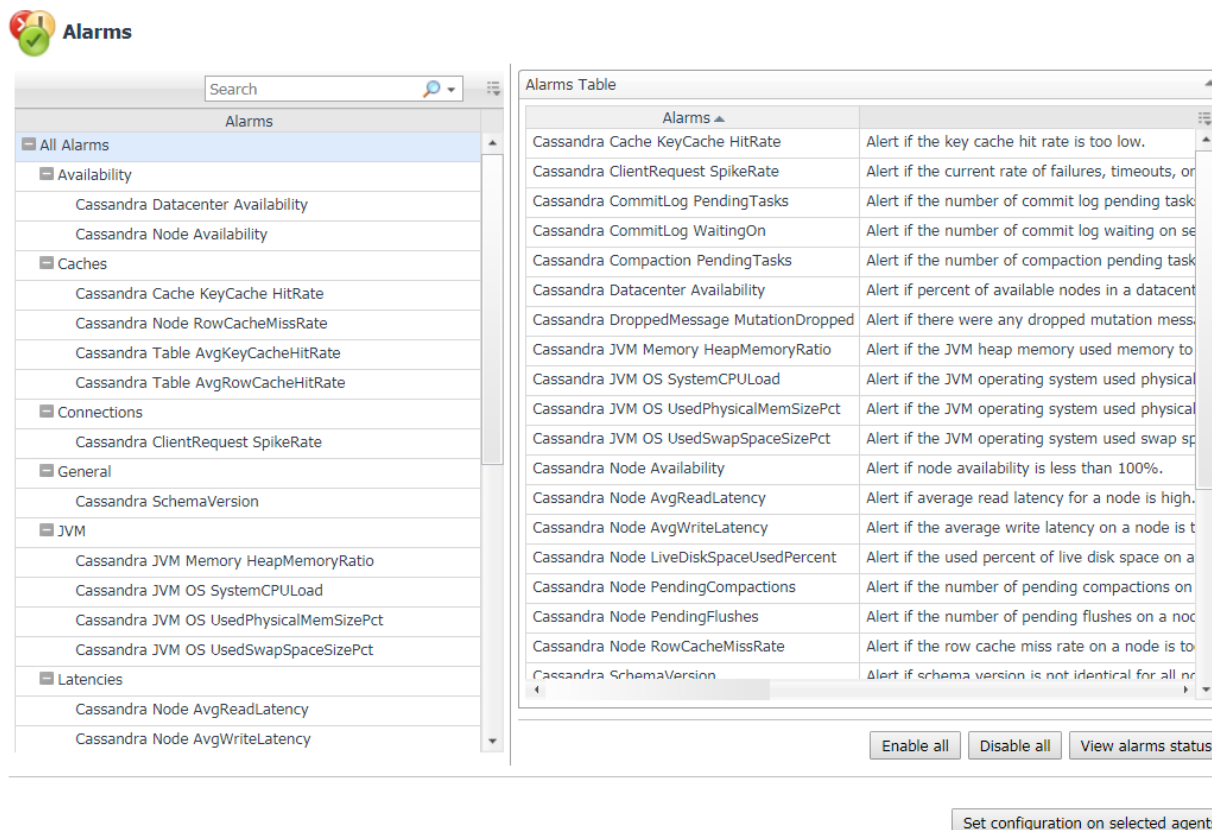
The Alarms view enables you to modify global settings and agent-specific settings for alarms.

To open the Alarms view:

1. Open the Administration dashboard as described in Opening the Databases Administration Dashboard.
2. Select the agents you wish to modify and do one of the following steps:
 - a. Select the Settings button and open the Administration dashboard, then click Alarms.
 - b. Select the 'Configure Alarm' button.
3. From the Alarms view, you can complete the following tasks:
 - a. [Modifying Alarm Settings](#)
 - b. [Reviewing Rule Definitions](#)
 - c. [Cloning Agent Settings](#)

Modifying Alarm Settings

You can customize how the alarms generated by the default rules are triggered and displayed in the Alarm view. Changes to alarm settings will apply to all selected agents, though thresholds can be customized by individual agent.



The screenshot shows the 'Alarms' interface. On the left, there is a sidebar with a search bar and a list of alarm categories: All Alarms, Availability, Caches, Connections, General, and JVM. Each category has a list of specific alarms. On the right, there is a table titled 'Alarms Table' with two columns: 'Alarms' and 'Alert'. The table lists various alarms and their corresponding alert messages. At the bottom of the table, there are buttons for 'Enable all', 'Disable all', and 'View alarms status'. Below the table, there is a button labeled 'Set configuration on selected agents'.

Alarms	Alert
Cassandra Cache KeyCache HitRate	Alert if the key cache hit rate is too low.
Cassandra ClientRequest SpikeRate	Alert if the current rate of failures, timeouts, or
Cassandra CommitLog PendingTasks	Alert if the number of commit log pending task
Cassandra CommitLog WaitingOn	Alert if the number of commit log waiting on se
Cassandra Compaction PendingTasks	Alert if the number of compaction pending task
Cassandra Datacenter Availability	Alert if percent of available nodes in a datacent
Cassandra DroppedMessage MutationDropped	Alert if there were any dropped mutation mess
Cassandra JVM Memory HeapMemoryRatio	Alert if the JVM heap memory used memory to
Cassandra JVM OS SystemCPULoad	Alert if the JVM operating system used physical
Cassandra JVM OS UsedPhysicalMemSizePct	Alert if the JVM operating system used physical
Cassandra JVM OS UsedSwapSpaceSizePct	Alert if the JVM operating system used swap sp
Cassandra Node Availability	Alert if node availability is less than 100%.
Cassandra Node AvgReadLatency	Alert if average read latency for a node is high.
Cassandra Node AvgWriteLatency	Alert if the average write latency on a node is t
Cassandra Node LiveDiskSpaceUsedPercent	Alert if the used percent of live disk space on a
Cassandra Node PendingCompactions	Alert if the number of pending compactions on
Cassandra Node PendingFlushes	Alert if the number of pending flushes on a noc
Cassandra Node RowCacheMissRate	Alert if the row cache miss rate on a node is to
Cassandra SchemaVersion	Alert if schema version is not identical for all n

The Alarms list controls the contents displayed to the right and the tasks that are available.

- **All Alarms** – Displays all rules with configured alarms and indicates whether alarms are enabled. In this view, you can enable or disable alarms for all the rules at once. You can also set email notifications and define mail server settings.
- **Category of rules** – Displays a set of related rules with configured alarms. In this view, you can enable or disable alarms and also set email notifications for the category of rules.
- **Rule name** – Displays the alarm status for the selected rule. If the rule has multiple severity levels, displays the threshold configured for each severity level. In this view, you can enable or disable the alarm, edit the alarm text, and edit severity levels and their thresholds. You can also set email notifications for the alarm.

You can complete the following tasks:

- [Enabling or disabling alarms for selected agents](#)
- [Modifying alarm threshold values](#)
- [Editing the text of the alarm message](#)

Your changes are saved separately and applied over the default rules. This protects you from software upgrades that may change the underlying default rules.

Enabling or disabling alarms for selected agents

You can override the global alarm sensitivity level setting for the selected agents. You can enable or disable alarms for all rules, a category of rules, or an individual rule.

To see descriptions of the rules, follow the steps described in [Reviewing Rule Definitions](#).

To enable or disable alarms:

1. Navigate to the Alarms view.
2. Decide on the scope for the change: all alarms, a category of rules, or a selected rule.
3. Complete the steps for the selected scope:

Scope	Procedure
All alarms	Click All Alarms . In the Alarms Settings tab, click either Enable all or Disable all .
Category of rules	Click a category. Click either Enable all or Disable all .
Selected rule	Click the rule. In the Alarms Settings tab, click the link that displays the alarm status. Select Enabled or Disabled from the list and click Set .

Modifying alarm threshold values

You can and should modify the thresholds associated with alarms to better suit your environment. If you find that alarms are firing for conditions that you consider to be acceptable, you can change the threshold values that trigger the alarm. You can also enable or disable severity levels to better suit your environment.

When a rule has severity levels, a Threshold section appears in the Alarm Settings tab showing the severity levels and bounds by agent. The threshold values correspond to the lower bounds shown in this table. Many rules do not have severity levels and thresholds.

When editing thresholds, ensure that the new values make sense in context with the other threshold values. For most metrics, threshold values are set so that Warning < Critical < Fatal. However, in metrics where normal performance has a higher value, such as DBSS - Buffer Cache Hit Rate, the threshold values are reversed: Warning > Critical > Fatal.

To change severity levels and thresholds:

1. Navigate to the Alarms view.
2. Click the multiple-severity rule that you want to edit.
3. Click the **Alarms Settings tab**.
4. In the Threshold section, review the defined severity levels and existing threshold bounds for all target agents.
5. Modify the severity levels for one or more agents by following one of the following procedures:

Task	Procedure
Edit severity levels and set threshold values for all agents.	Click Enhance alarm . Select the check boxes for the severity levels you want enabled and set the threshold values. Click Set .
Change the threshold values for one agent.	Click Edit beside the agent name. Set the new threshold values and click Set .
Copy the changes made to one agent's threshold values to all other agents.	Click Edit beside the agent name that has the values you want to copy. Select Set for all agents in table and click Set .

Editing the text of the alarm message

For individual rules, you can change the message displayed when an alarm fires. You cannot add or remove the variables used in the message. This is a global setting that affects all agents.

To change the alarm message:

1. In the Alarms view, click the **Settings** tab.
2. Select a rule.
3. Click the **Alarm Settings** tab.
4. Click **Enhance alarm**. A Customize <rule> dialog box opens.
5. In the Message box, edit the message text. To restore the default message, click **Reset message**.
6. Click **Set**.

Reviewing Rule Definitions

If you want to review the conditions of a rule, open the rule in the Rule Management dashboard.

IMPORTANT: Avoid editing rules in the Rule Management dashboard unless you are creating your own rules or copies. These rules may be modified during regular software updates and your edits will be lost.

You can create user-defined rules from the Rule Management dashboard. If you want to modify a rule, we recommend copying the rule and creating a user-defined rule. User-defined rules need to be managed from the Rule Management dashboard; these rules are not displayed in the Alarms view of the Databases Administration dashboard. For help creating rules, open the online help from the Rule Management dashboard.

To open the Rule Management dashboard:

1. On the navigation panel, under **Homes**, click **Administration**.
2. In the Administration dashboard, click **Rules**.
3. Type **Redshift** in the Search field to see the list of predefined rules for Redshift databases. The Redshift rules are displayed. From here, you can review threshold values, alarm counts, and descriptions.
4. To see the full rule definition, click a rule and then click **View and Edit**.
5. In the Rule Detail dialog box, click **Rule Editor**.
6. When you are done with your review, click Rule Management in the bread crumbs to return to the dialog box.
7. Click **Cancel** to avoid changing the rule unintentionally.

Cloning Agent Settings

You may want an agent to have the same settings as another agent. For example, if you add new agents, you may want them to use the same settings as an existing agent. In this case, you can clone the settings from one agent to other agents. This process does not link the agents; in the future if you update the source agent, you also need to update the target agents.

This procedure walks you through selecting the source agent from the Databases dashboard. However, you can also open the Administration dashboard with multiple agents selected. In this case, you select the source agent in Clone Alarm-related Settings to Other Agents dialog box.

To clone alarm-related settings:

1. On the Databases dashboard, select the check box for the agent with the settings you want to clone.
2. Click Settings and then Administration.
3. In the Administration dashboard, click Alarms.
4. Click Set configuration on selected agents. The Clone rule settings across agents dialog box opens.
5. In the Select the source agent drop-down list, you should see the agent you selected.
6. In the Select the target agents table, select the check boxes for agents that should inherit settings from the source agent.
7. Click Apply.
8. When prompted for confirmation, click Yes.

Configuring Email Notifications

We recommend that you set email notifications for the alarms you are most interested in tracking closely. For example, you may want to be notified by email of any Critical or Fatal situation. Or you may want to be informed whenever a key metric is no longer operating within acceptable boundaries.

You can set up email notifications that are generated when an alarm fires and/or on a defined schedule, as described in the following topics:

- [Configuring an email server](#)
- [Defining Default Email settings](#)
- [Enabling or disabling email notifications](#)
- [Defining email notifications, recipients, and messages](#)
- [Defining variables to contain email recipients](#)

Configuring an email server

You need to define the global mail server variables (connection details) to be used for sending email notifications.

The setting of the email should be configured in Foglight Administration > Email configuration.

Defining Default Email settings

You can define a default email address to be used by every new agent created in the future, by selecting the Default email button when configuring email notification.

The Email addresses entered are applied to all monitored agents not only for the agents that were selected to enter the Alarm administration.

Enabling or disabling email notifications

You can enable or disable email notifications for all alarms, a category of alarms, or a selected rule. Email notifications are sent only if all the following conditions are met:

- The alarm email notification setting is enabled for the affected rule.
- The alarm is triggered by changes in the monitored environment.
- Alarm notification is enabled at the triggered severity level. See Defining email notifications, recipients, and messages.

To enable or disable email notifications:

1. In the Alarms view, click the Settings tab.
2. Decide on the scope for the change: all alarms, a category of rules, or a selected rule.
3. Complete the steps for the selected scope:
 - All alarms - Click All Alarms. Click the Define Email Settings button. Select either Enabled or Disabled from the Alarms notification status list. Click Set.

- Category of rules - Click a category. Click the Define Email Settings button. Select either Enabled or Disabled from the Alarms notification status list. Click Set.
- Selected rule - Click a rule. In the Alarms Settings tab, click the Define Email Settings tab. Click the link that displays the alarm notification status. Select Enabled or Disabled and click Set.

Defining email notifications, recipients, and messages

You control who receives email messages, the subject line, and some text in the body of the email. The body of the email always contains information about the alarm. This information is not editable. You can also control whether an email is sent based on severity levels. You can set different distribution lists for different rules and different severity levels, or set the same notification policy for all rules.

To configure email notifications:

1. In the Alarms view, click the Settings tab.
2. Decide on the scope for the change: all alarms, a category of rules, or a selected rule.
3. Complete the steps for the selected scope:
 - All alarms - Click All Alarms. Click the Define Email Settings button.
 - Category of rules - Click a category. Click the Define Email Settings button.
 - Selected rule - Click a rule. Click the Email Notification Settings tab.
4. If you selected All Alarms or a category, in the Email Notification Settings dialog box, do one of the following:
 - To change the severity levels that warrant an email notification, from the Messages will be enabled for severities box, select the desired levels of severity.
 - To configure the same email recipients and message for all severity levels, click Configure mail recipients for all Severities and then click All severities.
 - To configure different email recipients and messages for each of the severity levels, click Configure mail recipients for the following options and then click a severity level.
5. In the Message Settings dialog box, configure the email recipients and message. Note that you can use registry variables in place of email addresses. Type the variable name between two hash (#) symbols, for example: #EmailTeamName#. For more information, see Defining variables to contain email recipients.
 - To — Type the addresses of the people who need to take action when this alarm triggers.
 - CC — Type the addresses of the people who want to be notified when the alarm triggers.
 - Subject — Optional. Edit the text of the subject line to better suit your environment. Avoid editing the variables, which are identified with the @ symbol.
 - Body Prefix — Optional. Add text that should appear above the alarm information in the body of the email.
6. Click Set to save the message configuration and close the dialog box.
7. If the Edit Notification Settings dialog box is open, click Set.

Defining variables to contain email recipients

You can create registry variables that contain one or more email addresses, and use these registry variables when defining email notifications. This procedure describes how to create a registry value.

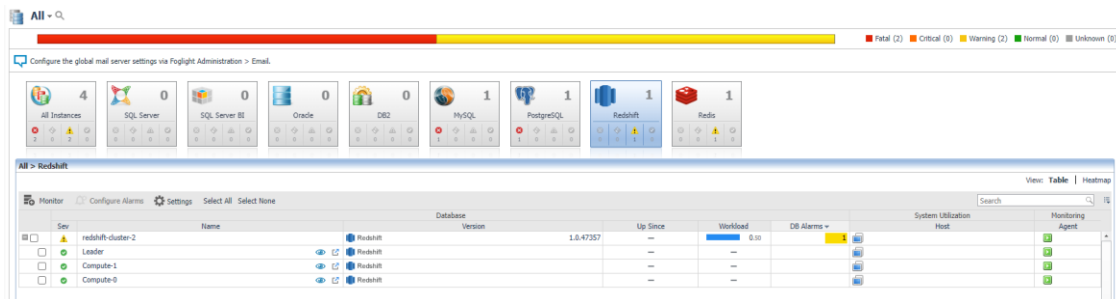
To create a registry variable:

1. On the navigation panel, under Dashboards, click Administration > Rules & Notifications > Manage Registry Variables.
2. Click Add. The New Registry Variable Wizard opens.
3. Select the registry variable type String, and click Next.
4. In the Name field, enter a name, for example: EmailTeamName
5. Click Next.
6. Select Static Value.
7. In the Enter desired value box, enter one or more email addresses (separated by commas).
8. Click Finish.

Dashboards

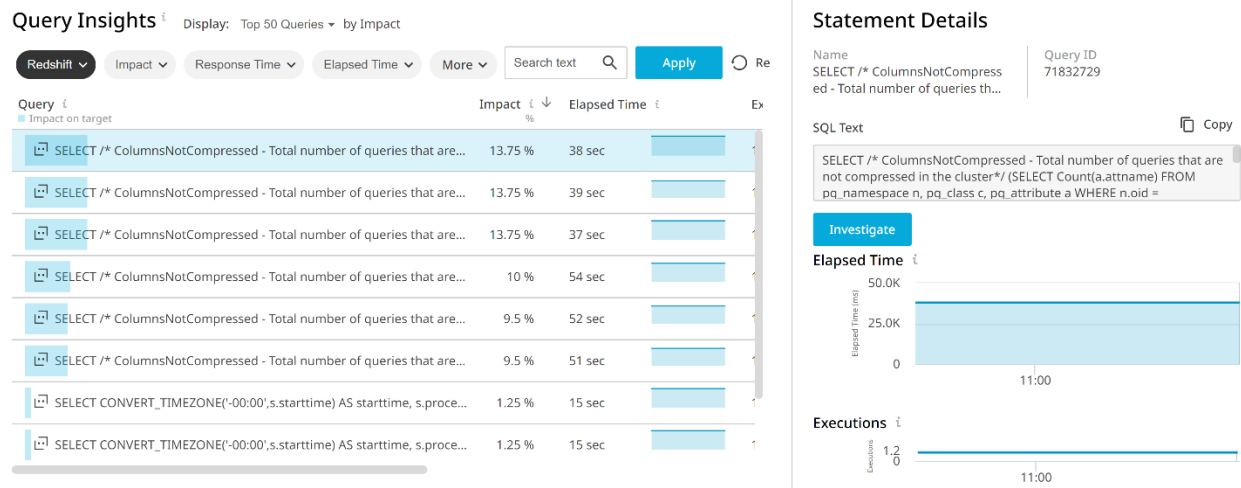
Databases

Foglight for Redshift is incorporated into the Databases dashboard along with many other monitored database types in your environment. Like other products, the list of databases can be filtered by type and severity level and includes basic information like health status, alarms, workload, and agent status. Because Redshift is a managed cloud service, certain information like host metrics are inapplicable. The cluster can be expanded to show individual nodes. Clicking a name will drill down to the Cluster Overview or Nodes pages.



Query Insights

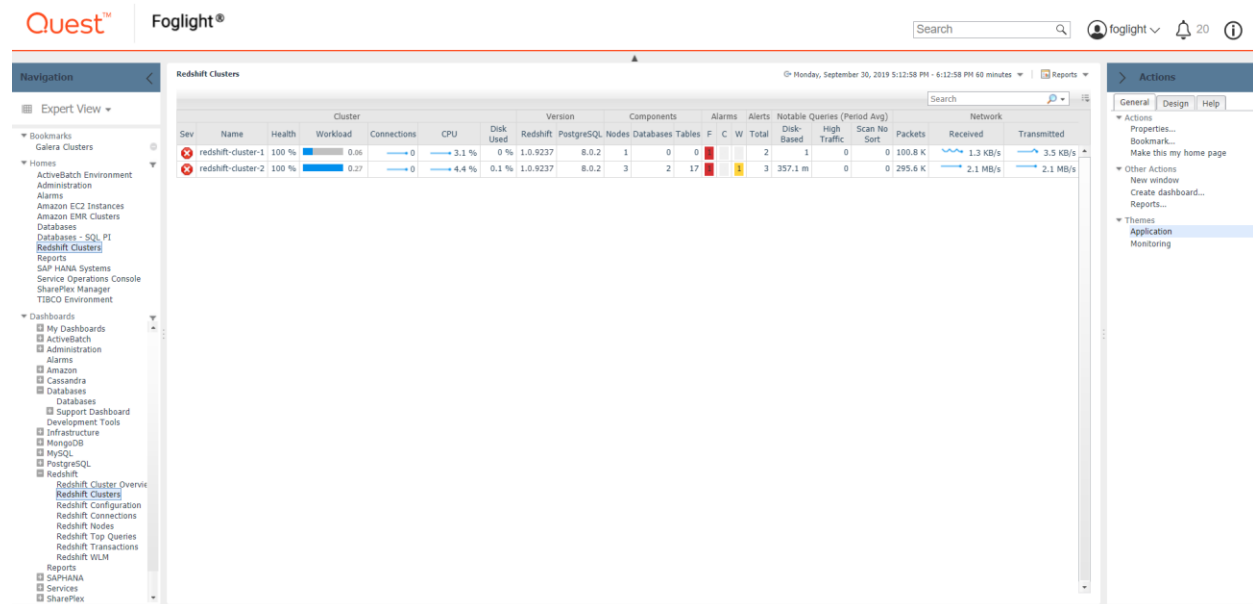
The Query Insights dashboard shows the top queries across all monitored databases along with their key metrics and estimated impact percent on the instance (statement workload as a percentage of total workload for the instance in the selected time range). Queries can be filtered and sorted by domain, SQL text, and any of the related metrics. Selecting a query will update the details panel on the right a focus on that query and clicking the Investigate button will go to the relevant page in the cartridge dashboards with all collected information for that query.



Redshift Clusters

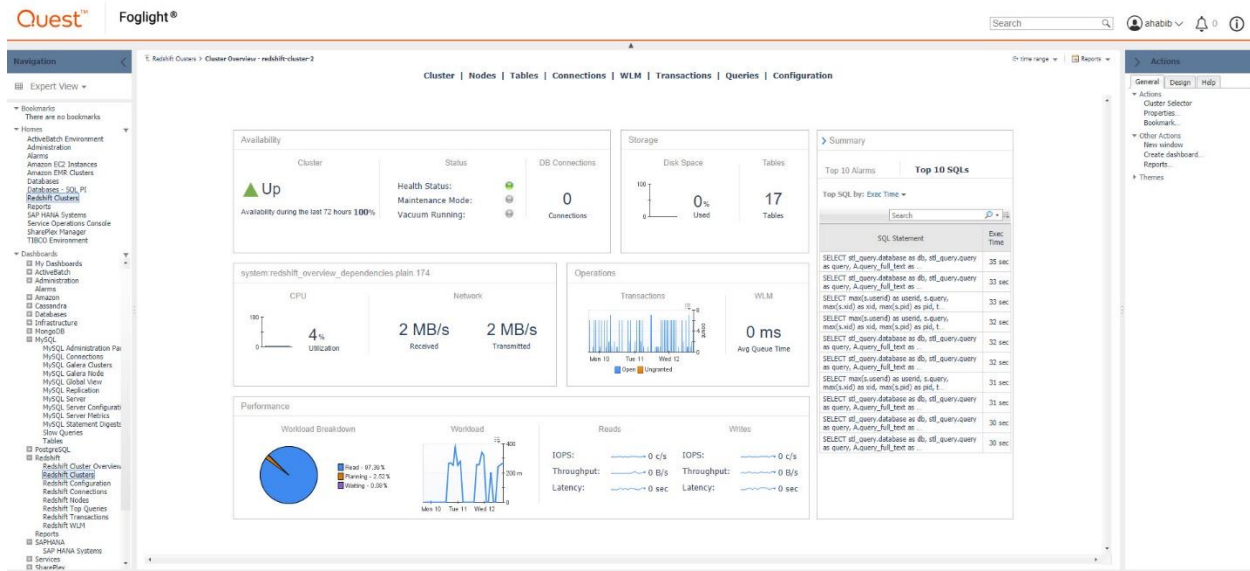
This dashboard lists all monitored Redshift Clusters and contains high-level information on cluster structure, nodes, health status, notable queries and key metrics. The workload metric shows the

amount of work the server is doing to respond to user queries and can be used to compare workload against other clusters. Selecting a cluster will navigate to the Cluster Overview page. In the cluster table, when the value in the Nodes column is clicked, it will navigate to the Nodes page and similarly, when the Databases or Tables value is clicked, it will navigate to the Tables page. Other column values will show a time plot when hovered or clicked on.



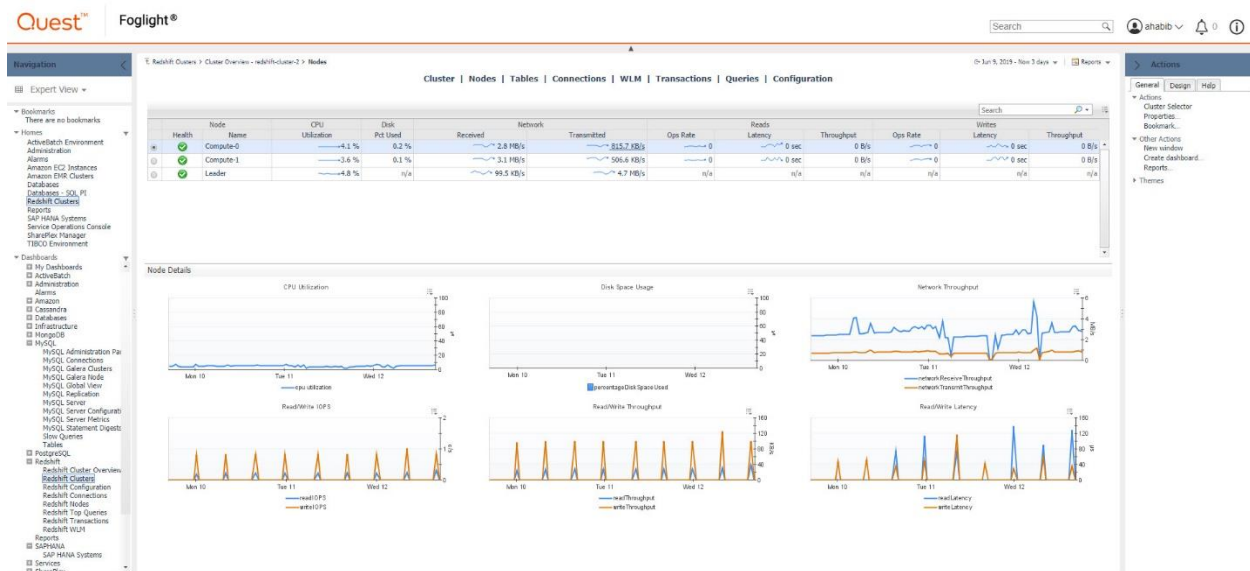
Cluster Overview

This page displays the cluster overview with availability, performance, storage metrics and the top 10 alarms and SQL statements. The Workload Breakdown pie chart can be clicked or hovered over for a detailed chart of time spent in different query execution stages for the selected time range. The SQL statements can be ordered by the metrics in the dropdown box.



Nodes

This page lists all nodes in the selected cluster. Selecting a node will update the bottom section of the page and display performance metrics of the node showing CPU Utilization, Disk Space usage, Network Throughput, Read/Write IOPS, Throughput and Latency. When the Health icon is clicked or hovered on, it displays a list of alarms for that node. All other column values will show a time plot when hovered or clicked on.



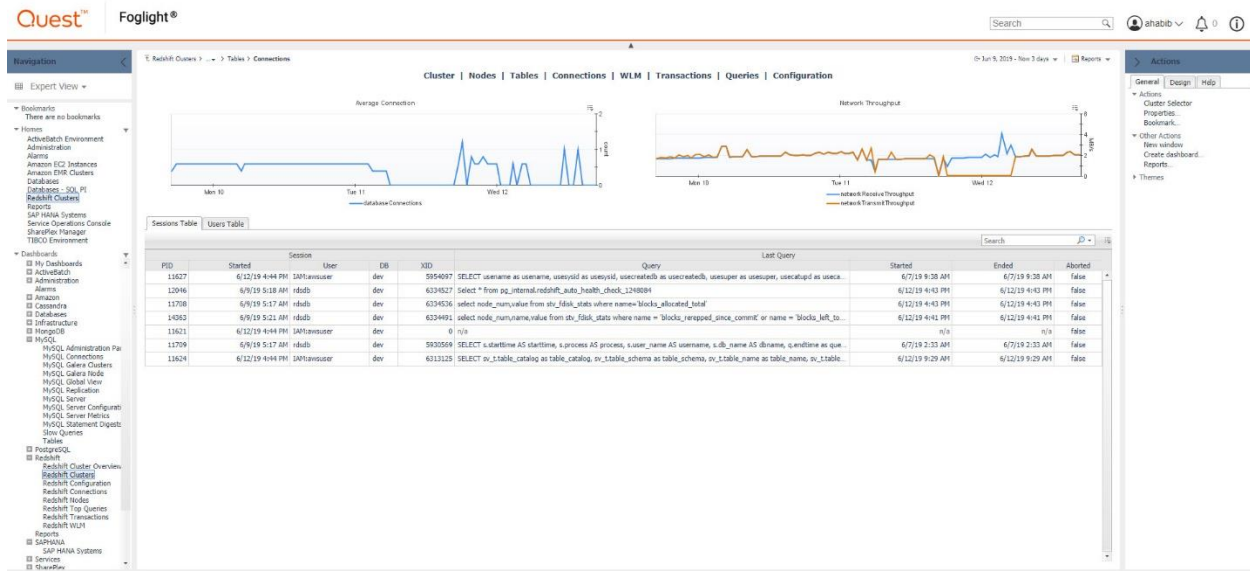
Tables

This page lists all tables in the selected cluster with information on the size, number of rows, percentage of space used by table and information on sort keys and key distribution. When the health icon is clicked or hovered on, it displays a list of alarms for that table. All other metrics will show a time plot when hovered or clicked on. To filter the list of tables by database, click the select database button on the top left of the table and select one or more databases for which you wish to view the tables.

Health	Database	Name	Rows	Size	Pct Used	Unsorted	Skew	Skew Rows
✓	workload0	event	8.8 K	72 MB	0 %	0 %	1 %	1 %
✓	dev	redshift_auto_health_check_1240094	4 K	16 MB	0 %	n/a	n/a	n/a
✓	dev	shoes	4	15 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable0	62	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable1	89	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable2	72	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable3	148	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable4	127	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable5	99	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable6	86	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable7	418	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable8	22	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable9	60	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable10	147	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable11	105	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable12	164	20 MB	0 %	n/a	n/a	n/a
✓	workload0	testtable13	201	20 MB	0 %	n/a	n/a	n/a

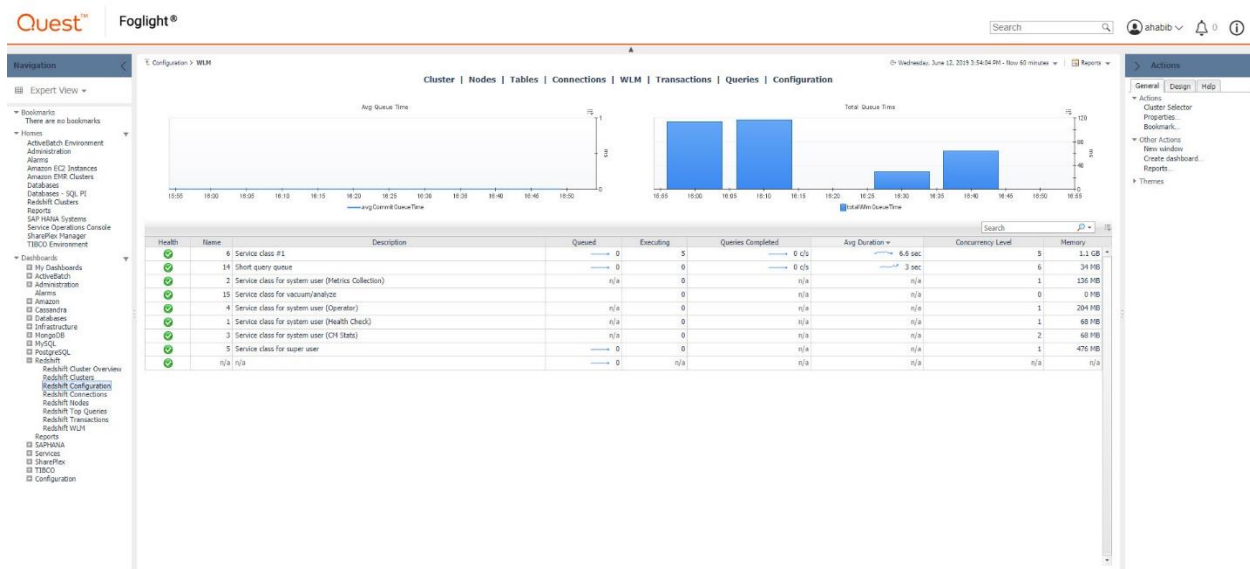
Connections

The Connections page displays the active user sessions and user information. The Sessions Table displays the session information and the associated query information with the user. The Users Table displays a list of database users and their associated privileges. The time plots above display the average number of Database Connections and Network Throughput.



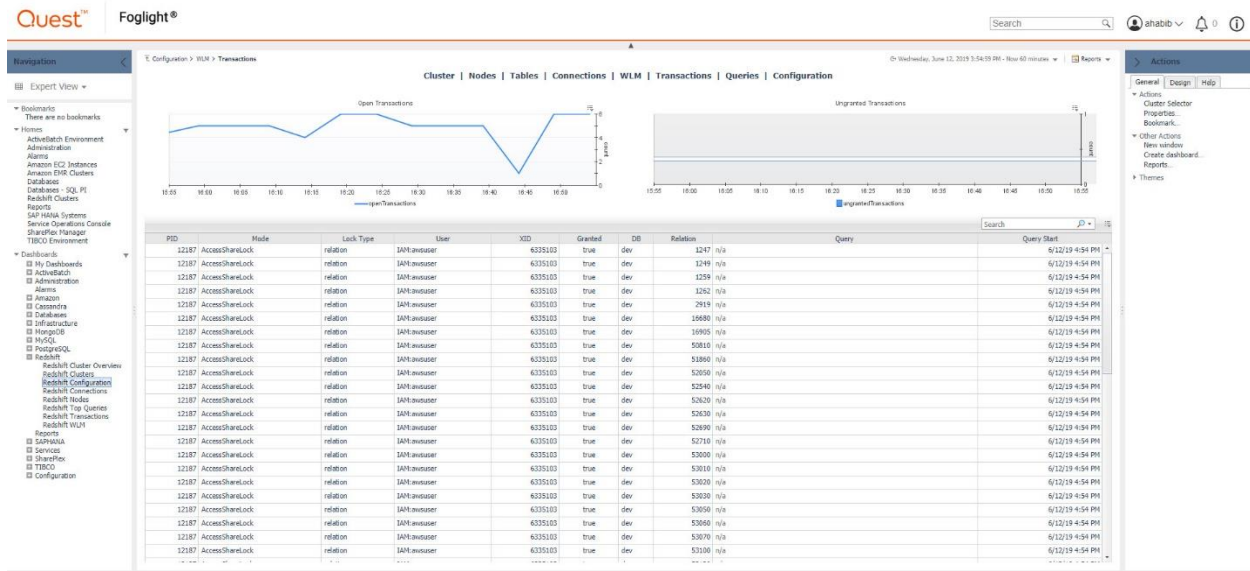
WLM

This page displays the key configurations and the current state of the service classes for WLM. When the health icon is clicked or hovered on, it displays a list of alarms for that queue. All other metrics will show a time plot when hovered or clicked on. There is a time plot and a time bar above which displays Average Commit Queue Time and Total WLM Queue Time respectively.



Transactions

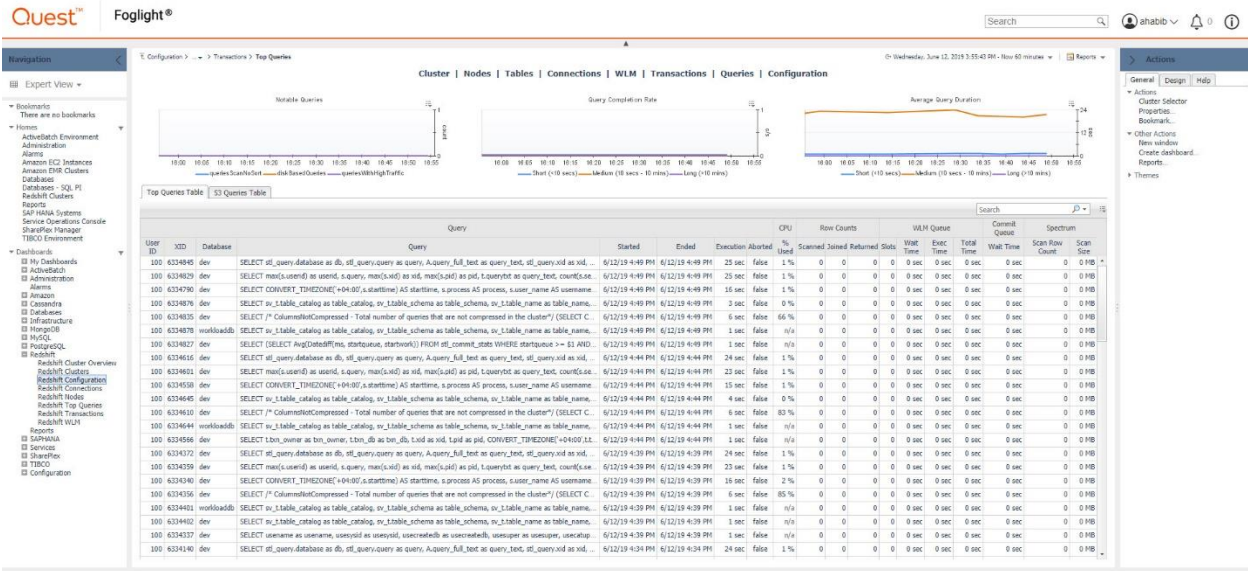
This page displays vital information about transactions that currently hold locks on tables in the database. The time plots above display Open Transactions – the number of open transactions during each sample time – and Ungranted Transactions – transactions where the lock has been pending.



Queries

The Top Queries table lists top queries and associated information that were executed in the selected page time range. Top queries for each sample period are determined by configuration set in the Agent Properties' Statement Tracking section. When hovered or clicked on, the query value displays the entire query text.

The S3 Queries table displays Redshift Spectrum queries that have been run on the system in the selected page time range. The S3 queries can be limited by enabling S3 query tracking and providing a limit in # of Top Queries in the Agent Properties' Statement Tracking section. When hovered or clicked on, the query value displays the entire query text.



Configuration

This page displays the runtime configuration parameters of the cluster with name, current value, descriptions and other properties.

The screenshot shows the Foglight interface for a Redshift cluster. The 'Configuration' tab is active, displaying a table of configuration parameters. The table columns include Name, Current Setting, Initial Setting, Short Desc, Property, Var Type, Source, Changed, Last Modified, History, and Total. The parameters are sorted by name, with the top parameter being 'query_group'.

Name	Current Setting	Initial Setting	Short Desc	Property	Var Type	Source	Changed	Last Modified	History	Total
query_group	default	default			string	configuration file	No	04/26/19 20:40	0	0
search_path	schema, public	schema, public	Sets the schema search order for names that are not schema-qualified.		string	configuration file	No	04/26/19 20:40	0	0
datestyle	ISO, MDY	ISO, MDY	Sets the display format for date and time values.		string	configuration file	No	04/26/19 20:40	0	0
analyze_percent	10	10	percentage of rows changed to trigger full analysis		real	default	No	04/26/19 20:40	0	0
statement_timeout	0	0	Sets the maximum allowed duration (in milliseconds) of any statement.		integer	configuration file	No	04/26/19 20:40	0	0
extra_float_digits	2	2	Sets the number of digits displayed for floating-point values.		integer	client	No	04/26/19 20:40	0	0
vlim_query_slot_count	1	1	number of vlim query slots to be acquired by queries		integer	default	No	04/26/19 20:40	0	0

Rules

Redshift Cluster Availability

Alert if the agent is unable to connect to the cluster.

Redshift CPU Utilization

Alert if a cluster is utilizing most of the available CPU.

Redshift Node Disk Space

Alert if the disk space usage of a node is approaching capacity.

Redshift Table Out of Date Statistics

Alert if a table's statistics are out of date.

Redshift Table Skew Value

Alert if table's skew value exceeds recommended limit.

Redshift Table Unsorted Percent

Alert if a table's unsorted percent is exceeding limit.

Redshift WLM Query Slots

Alert if the WLM query slot count exceeds the recommended limit.

Redshift WLM Queue Bottleneck

Alert if a WLM's Queue length is approaching capacity.

Reports

Redshift Cluster Summary

Cluster summary of a Redshift cluster including health, workload, CPU utilization, disk space, read/write latency and WLM average commit queue time.

Redshift S3 Queries

Displays the Top S3 Queries for a Redshift cluster for the provided time range.

Redshift Tables

Displays Top Tables for a Redshift cluster, sorted by a provided criterion.

Redshift Top Queries

Displays the Top Queries for a Redshift cluster for the provided time range.