

Foglight for Cassandra 6.0.1

Cartridge Guide

© 2021 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 | 6 |
| Description | 6 |
| Business Challenge..... | 6 |
| Key Features | 6 |
| Foglight for Cassandra Requirements | 7 |
| Installing and Configuring Agents | 8 |
| Cassandra Pre-Configuration..... | 9 |
| Database User Setup..... | 9 |
| JMX Setup..... | 10 |
| Configuring an SSL Connection..... | 11 |
| Cartridge Installation | 13 |
| Creating and Configuring Agents | 14 |
| Using the Agent Installer Wizard..... | 15 |
| Using the Agent Status Dashboard..... | 16 |
| Agent Properties..... | 17 |
| DB Connection | 17 |
| JMX Connections | 18 |
| Collection Periods | 18 |
| Options | 18 |
| Roles | 21 |
| Upgrading the Agent | 22 |
| Removing Monitored Databases..... | 23 |
| Administration | 24 |
| Opening the Databases Administration Dashboard | 24 |
| Reviewing the Administration Settings..... | 24 |
| Customizing Alarms for Foglight for Cassandra Rules | 25 |
| Introducing the Alarms View | 25 |

| | |
|---|-----------|
| Modifying Alarm Settings | 26 |
| Reviewing Rule Definitions | 29 |
| Cloning Agent Settings..... | 30 |
| Configuring Email Notifications | 31 |
| Dashboards | 34 |
| Databases | 34 |
| Cassandra Clusters..... | 35 |
| Cluster Keyspaces | 36 |
| Cluster Tables | 37 |
| Traces Sample | 38 |
| Node Overview..... | 39 |
| Node Keyspaces | 39 |
| Node Tables | 40 |
| Node Table..... | 40 |
| Client Requests | 41 |
| Connections | 41 |
| JVM | 42 |
| Cache | 43 |
| Thread Pools..... | 44 |
| Slow Query Log..... | 44 |
| Rules | 45 |
| Cassandra Cache KeyCache HitRate..... | 45 |
| Cassandra ClientRequest SpikeRate | 45 |
| Cassandra CommitLog PendingTasks | 45 |
| Cassandra CommitLog WaitingOn | 45 |
| Cassandra Compaction PendingTasks..... | 45 |
| Cassandra Datacenter Availability | 45 |
| Cassandra DroppedMessage MutationDropped..... | 45 |
| Cassandra JVM Memory HeapMemoryRatio | 45 |
| Cassandra JVM OS SystemCPULoad..... | 45 |
| Cassandra JVM OS UsedPhysicalMemSizePct | 46 |

| | |
|---|-----------|
| Cassandra JVM OS UsedSwapSpaceSizePct | 46 |
| Cassandra Node Availability | 46 |
| Cassandra Node AvgReadLatency | 46 |
| Cassandra Node AvgWriteLatency | 46 |
| Cassandra Node LiveDiskSpaceUsedPercent | 46 |
| Cassandra Node PendingCompactions | 46 |
| Cassandra Node PendingFlushes | 46 |
| Cassandra Node RowCacheHitRate | 46 |
| Cassandra SchemaVersion | 47 |
| Cassandra Storage Exceptions..... | 47 |
| Cassandra Storage LoadVsSpace..... | 47 |
| Cassandra Storage TotalHints | 47 |
| Cassandra Storage TotalHintsInProgress | 47 |
| Cassandra Table AvgKeyCacheHitRate..... | 47 |
| Cassandra Table AvgReadLatency | 47 |
| Cassandra Table AvgRowCacheHitRate | 47 |
| Cassandra ThreadPools PendingTasks | 47 |
| Cassandra ThreadPools TotalBlockedTasks..... | 48 |
| Reports..... | 49 |
| Cassandra Cluster Summary | 49 |
| Cassandra Cluster Tables..... | 49 |
| Cassandra Node Executive Summary | 49 |
| Cassandra Node Health Check..... | 49 |
| Cassandra Storage Report | 49 |

Introduction

Description

Complex applications that require storage of a massive amount of data and a flexible data structure require a new form of database. Cassandra database offers scalability and high availability without compromising performance on commodity hardware or cloud infrastructure. This makes it an ideal platform for mission-critical data and for replicating across multiple datacenters. Cassandra is a best-in-class solution, providing low latency for users and peace of mind, with an architectural rigidity designed to survive regional outages.

Business Challenge

Modern NoSQL databases like Cassandra are designed to support massive data processing and provide an equally large storage capability, but hosting Cassandra clusters with a large number of nodes can greatly increase the complexity of your data infrastructure. Understanding the performance of your Cassandra clusters is critical for diagnosing issues and planning capacity.

Foglight for Cassandra provides comprehensive performance monitoring and administration for all nodes in a cluster from a centralized console. You can collect statistical data from all JVMs in a cluster and key performance metrics like memory utilization statistics, task statistics of thread pools, storage statistics, CPU usage, operational performance, latency, and bottlenecking.

Key Features

Foglight for Cassandra provides a consolidated view of all monitored **Cassandra Clusters**, containing information on cluster structure, nodes, health status, and other key metrics. Derived calculations provide insight into the overall cluster workload taking into account activity on each node in the cluster.

View **Cluster Health** at a glance, displaying a topology of the cluster structure or a list of nodes with relevant performance and availability metrics. **Tables** can be viewed by cluster or namespace, aggregating associated metric information across nodes. **Node** Information is granular and includes Node Health and Alerts, Availability, Workload, Statements, Connections and Operations.

Understand every component of Cassandra performance. Granular **performance metrics** such as Reads/Writes, Caches, Bloom Filter, Buffer Pool, Memtables, and Commit Log are collected and stored for historical and trend analysis. JVM information is also collected and analyzed to help ensure optimal configuration and performance.

Understand response times with a breakdown of **Client Requests** to the node by type, showing metric histories for average latency, request counts, and request errors, broken down by error type. Understand the performance and availability between all nodes in your cluster. Foglight for Cassandra monitors all **Connections** including dropped messages, timeouts, and pending messages.

Alerts can be generated based on preset rules and thresholds as well as dynamically, based on situational anomalies. Alerting has been designed to help ensure the performance and availability of your Cassandra environment at all times and to keep it performing optimally.

Foglight for Cassandra Requirements

Foglight for Cassandra is compatible with **Apache Cassandra versions 2.1+ and DSE versions 4.8+**. However, some data may not be available in earlier versions. The following list itemizes which features become available at which versions.

v3.0:

- Traces include client and command.

- Buffer pool JMX metrics available.

- Storage service JMX metrics available.

- Various changes in JMX table metrics.

Foglight for Cassandra can be installed on **FMS 6.0.0+** and agents require **FglAM 6.0.0+**.

Installing and Configuring Agents

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

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

Cassandra Pre-Configuration

In order to allow full monitoring of Cassandra, the agent will require a user with sufficient privilege to execute system queries. For JMX monitoring, remote JMX connections must also be enabled on each node.

Database User Setup

Using a superuser role or a role with user creation permissions, connect via `cqlsh` to a node on the cluster and create a new user for the Foglight Agent as follows:

Cassandra >= 2.2

For versions of Cassandra >= 2.2 or DSE >= 5.0, create a user as follows:

```
CREATE ROLE IF NOT EXISTS <username> WITH PASSWORD = '<your-password>' AND LOGIN = true
AND SUPERUSER = false;
```

There are then two ways to set permissions. The simpler way is to use the following single permission:

```
GRANT SELECT ON ALL KEYSPACES to <username>;
```

It is also possible to use more specific permissions:

```
GRANT SELECT ON TABLE system.local to <username>;
GRANT SELECT ON TABLE system.peers to <username>;
GRANT SELECT ON TABLE system_traces.events to <username>;
GRANT SELECT ON TABLE system_traces.sessions to <username>;
GRANT SELECT ON TABLE system_auth.roles to <username>;
```

Using specific permissions, there is one version-dependent permission. For versions of Cassandra >= 3.0:

```
GRANT SELECT ON TABLE system_schema.tables to <username>;
```

For versions of Cassandra >= 2.2 but less than 3.0:

```
GRANT SELECT ON TABLE system.schema_columnfamilies to <username>;
```

Using the more specific permissions on any version >=2.2 use the following if running DSE:

```
GRANT SELECT ON TABLE dse_perf.node_slow_log to <username>;
```

Cassandra < 2.2

For versions of Cassandra < 2.2 or DSE <= 4.8:

```
CREATE USER IF NOT EXISTS <username> WITH PASSWORD '<your-password>' NOSUPERUSER;
GRANT SELECT ON ALL KEYSPACES TO <username>;
```

JMX Setup

Enable Remote JMX Connections

Enable remote JMX authentication as per the instructions at either of the following links:

<https://docs.datastax.com/en/cassandra-oss/3.x/cassandra/configuration/secureJmxAuthentication.html>

<https://cassandra.apache.org/doc/latest/operating/security.html?#standard-jmx-auth>

JMX Remote Files

Only a single user needs to be added to the *jmxremote.password* and *jmxremote.access* files for Foglight monitoring and administration. The following is the full (template) content of a sample *jmxremote.password* file:

```
<your-JMX-username> <your-JMX-password>
```

The following is the full (template) content of a sample *jmxremote.access* file:

```
<your-JMX-username> readwrite
```

There are two JMX access permissions to choose between:

- `readonly` – For monitoring only.
- `readwrite` – For monitoring as well as administration, which is currently the ability to set the per-node trace probability.

Configuring an SSL Connection

The below instructions cover common steps used to configure a TLS/SSL connection from the Cassandra Agent client. A full treatment of TLS/SSL keys, certificates, and certificate authorities (CA) is beyond the scope of this document. The following instructions assume familiarity with TLS/SSL concepts and tools. Client and certificate authority certificates must be available prior to proceeding.

In order to use SSL, your Cassandra cluster must include SSL support and allow SSL connections. There are various configurations options for client connections. Refer to the Cassandra documentation and verify that the current Cassandra server configuration parameters support the desired authentication.

The Foglight agent, in its capacity as a database client, requires access to a private key, its signed certificate, and the signing CA's certificate. The client key and certificate must be imported into a keystore, and the CA certificate must be imported into a separate truststore.

One example method for generating a JKS keystore for use with Foglight utilizes openssl and keytool. Set the key and certificate filenames, alias name, and keystore password as appropriate.

```
openssl pkcs12 \  
    -export \  
    -in ${CERT_NAME}.crt \  
    -inkey ${CERT_NAME}.key \  
    -name $CERT_NAME \  
    -out temp-keystore.p12 \  
    -passout pass:${KEYPASS}  
  
keytool -importkeystore \  
    -srckeystore temp-keystore.p12 \  
    -srcstoretype PKCS12 \  
    -srcstorepass $KEYPASS \  
    -destkeystore keystore \  
    -deststoretype JKS \  
    -deststorepass $KEYPASS
```

Regardless of how the keystore is constructed, it must list the client certificate as a 'PrivateKeyEntry', indicating that it also contains the private key, not just the signed certificate.

Separately, import the CA certificate into a truststore:

```
keytool -importcert \  
    -keystore truststore \  
    -alias $CA_NAME \  
    -file ${CA_NAME}.crt \  
    -keypass $TRUSTPASS \  
    -storepass $TRUSTPASS \  
    -storetype JKS \  
    -noprompt
```

Next, edit the baseline.jvmargs.config file in the <FglAM-install-root>/state/default/config directory and add the following parameters with file paths and passwords appropriate for your system.

```
vmparameter.0 = "-Djavax.net.ssl.keyStore=/path/to/keystore";
```

```
vmparameter.1 = "-Djavax.net.ssl.keyStorePassword=changeit";  
vmparameter.2 = "-Djavax.net.ssl.trustStore=/path/to/truststore";  
vmparameter.3 = "-Djavax.net.ssl.trustStorePassword=changeit";
```


Escape any quotes with a backslash (\). On an Agent Manager installed on Windows, use forward slashes in the file paths, like so:

```
vmparameter.0 = "-Djavax.net.ssl.keyStore=\"C:/path/to/keystore\"";
```

Then, restart the FglAM and continue with the agent configuration, setting the “Use TLS/SSL?” option in the Agent Properties to true for the driver and/or JMX connections, depending on the Cassandra server configuration.

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 **CassandraAgent-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 Cassandra 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](#)
- [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

Foglight for Cassandra 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. Foglight for Cassandra allows running a wizard that provides a common entry point for adding database instances and then configuring these instances for monitoring.

To run the instance installation wizard:

1. On the navigation panel, click Homes > Databases.
2. Click the Cassandra 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 Cassandra 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 Cassandra 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 CassandraAgent 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 Cassandra agent.
4. Click the **Edit Properties** button.
5. Select **Modify the default properties for this agent**.
6. Edit the agent properties for the Cassandra agent instance:
 - [DB Connection](#)
 - [JMX Connections](#)
 - [Collection Periods](#)
 - [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 Cassandra 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.

| | |
|-------------------------------------|--|
| DB Connection | |
| IP or Hostname | <input type="text" value="ec2-11-111-111-111.compute-1.amazonaws.com"/> |
| Port | <input type="text" value="9042"/> |
| Username | <input type="text" value="foglightagent"/> |
| Password | <input type="password" value="••••••••"/> |
| Use TLS/SSL? | <input type="checkbox"/> |
| JMX Connections | |
| Node Connections | <input type="text" value="C_01"/> <input type="button" value="Edit"/> <input type="button" value="Clone"/> <input type="button" value="Delete"/> |
| Default Port | <input type="text" value="7199"/> |
| Default Username | <input type="text" value="cassandra"/> |
| Default Password | <input type="password" value="••••••••"/> |
| Use TLS/SSL? | <input type="checkbox"/> |
| Collection Periods | |
| Collection Periods (sec) | <input type="text" value="DefaultPeriods"/> <input type="button" value="Edit"/> <input type="button" value="Clone"/> <input type="button" value="Delete"/> |
| Options | |
| Host Aliases | <input type="text" value="HostAliasesList"/> <input type="button" value="Edit"/> <input type="button" value="Clone"/> <input type="button" value="Delete"/> |
| Contact Points | <input type="text" value="ContactPointList"/> <input type="button" value="Edit"/> <input type="button" value="Clone"/> <input type="button" value="Delete"/> |
| Max Traces/Queries | <input type="text" value="250"/> |
| Max Tables (0 for none, -1 for all) | <input type="text" value="-1"/> |

DB Connection

The agent requires a connection to the cluster in order to gather information about the cluster and data structure. The agent may be referred to other nodes than the one specified.

- **IP or Hostname** – Host where Cassandra node is running. Default is “localhost”. (e.g.<hostname> or <IP address>)
- **Port** – The CQL native transport port for the Cassandra node. Default is 9042.
- **Username** – User that can connect to the Cassandra node.
- **Password** – Password of the user that can connect to the Cassandra node.
- **Use TLS/SSL?** – Whether the database connection should be made over SSL.

JMX Connections

The agent requires JMX access to individual nodes in the cluster to gather most operational metrics.

- **Node Connections** – IP/Hostname for each node in the cluster to be monitored must be added to this list. Values for the columns Port, Username, and Password can be specified for each node if they differ. If no value is provided for these, values will be taken from the default fields under JMX Connections (see below). Check "Use as Host Alias?" to indicate that the value provided for "IP or Hostname" should be used as an alias for the discovered hostname for the respective node. Secondary Property Lists are global and can be shared between agents. To create a new property list for a different cluster, clone an existing list and then edit and save it as the selected list for that agent.
- **Default Port** – JMX listening port for the Cassandra node. This value will be used as the default value in the Node Connections list if no other value is provided. Default is 7199.
- **Default Username** – JMX User that can connect to the Cassandra node. This value will be used as the default value in the Node Connections list if no other value is provided.
- **Default Password** – Password of the JMX User that can connect to the Cassandra node. This value will be used as the default value in the Node Connections list if no other value is provided.
- **Use TLS/SSL?** – Whether the JMX connection should be made over SSL.

Collection Periods

The Collection Periods field list in the agent properties is used to set the sample frequencies. A collection can be turned off by setting its period to 0. The defaults are set based on the type of data being collected for relevancy.

- **Availability** – Controls the query interval for node availability collections.
- **Cluster and Node** – Controls the query interval for cluster and node collections over the database driver.
- **Keyspace and Table** – Controls the query interval for keyspace and table collections over the database driver.
- **JMX Cluster** – Controls the query interval for cluster metrics over JMX.
- **JMX Keyspace** – Controls the query interval for keyspace metrics over JMX. If Max Tables is set to zero keyspace metrics are not collected.
- **JMX Table** – Controls the query interval for table metrics over JMX. If Max Tables is set to a positive value, metrics controlled by this period are collected for the largest tables only. If Max Tables is equal to zero this period is not used. If Max Tables is set to -1 this period is used to collect metrics from all tables.
- **JMX Table Size** – Controls the query interval for table size metrics over JMX. If Max Tables is set to a positive value, metrics controlled by this period are collected for all tables for the purpose of determining the largest tables on the cluster. If Max Tables is less than or equal to zero this period is not used.
- **Slow Query** – Controls the query interval for collection of the DSE Slow Query log.
- **Trace** – Controls the query interval for collection of the trace log. Note that tracing must be enabled on the Cassandra instance(s) separately for data to be collected.
- **JVM** – Controls the query interval for the JVM collections over JMX.

Options

Host Aliases

Aliases can optionally be added for Cassandra nodes in case hostnames as discovered by the agent differ from those used for the Hosts monitor. This Secondary Property List is a pairing of the uniquely identifying Host ID (a UUID string) and the desired hostname alias. There are two ways to determine a

node's Host ID for use in the Aliases list. For a node that has already been monitored for some time, the Host ID can be found on the Nodes Table of the Cassandra Clusters dashboard (the Host ID column is hidden by default, so it will need to be made visible via the table's "Show columns" dialog). Host IDs can also be queried directly from a specific Cassandra instance with the following CQL:

```
SELECT host_id FROM system.local WHERE key='local';
```

Contact Points

Additional contact points for the database driver to assist with making an initial connection if the node specified by the main contact point is not available.

Max Traces/Queries

The maximum number of traces or slow query entries to collect per their respective collection periods. Limiting the number here prevents the FMS from being overwhelmed by data (particularly from traces) on very active instances.

Max Tables

The number of tables for which to collect full metrics. If set to 0, no keyspace or table metrics are collected. If set to the default of -1, full metrics are collected for all tables.

Roles

Two roles, Cassandra User and Cassandra Administrator, are installed with the cartridge. Viewing Cassandra dashboards requires that a user be assigned one of these or have the core Administrator role. The Set Trace Probability function on the Traces Sample dashboard requires the Cassandra Administrator role.

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

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 Cassandra dashboard.

Administration

Opening the Databases Administration Dashboard

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

NOTE: If you attempt to select instances of more than one type of database, such as a Cassandra 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 > Cassandra**.
2. Select the check boxes beside one or more Cassandra 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 Cassandra 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 Cassandra Rules

Many Foglight for Cassandra 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.

Alarms

Search

Alarms

- All Alarms
- Availability
 - Cassandra Datacenter Availability
 - Cassandra Node Availability
- Caches
 - Cassandra Cache KeyCache HitRate
 - Cassandra Node RowCacheMissRate
 - Cassandra Table AvgKeyCacheHitRate
 - Cassandra Table AvgRowCacheHitRate
- Connections
 - Cassandra ClientRequest SpikeRate
- General
 - Cassandra SchemaVersion
- JVM
 - Cassandra JVM Memory HeapMemoryRatio
 - Cassandra JVM OS SystemCPULoad
 - Cassandra JVM OS UsedPhysicalMemSizePct
 - Cassandra JVM OS UsedSwapSpaceSizePct
- Latencies
 - Cassandra Node AvgReadLatency
 - Cassandra Node AvgWriteLatency

| Alarms Table | |
|--|---|
| Alarms ▲ | |
| 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 |

Enable all Disable all View alarms status

Set configuration on selected agents

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 **Cassandra** in the Search field to see the list of predefined rules for Cassandra databases. The Cassandra 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.

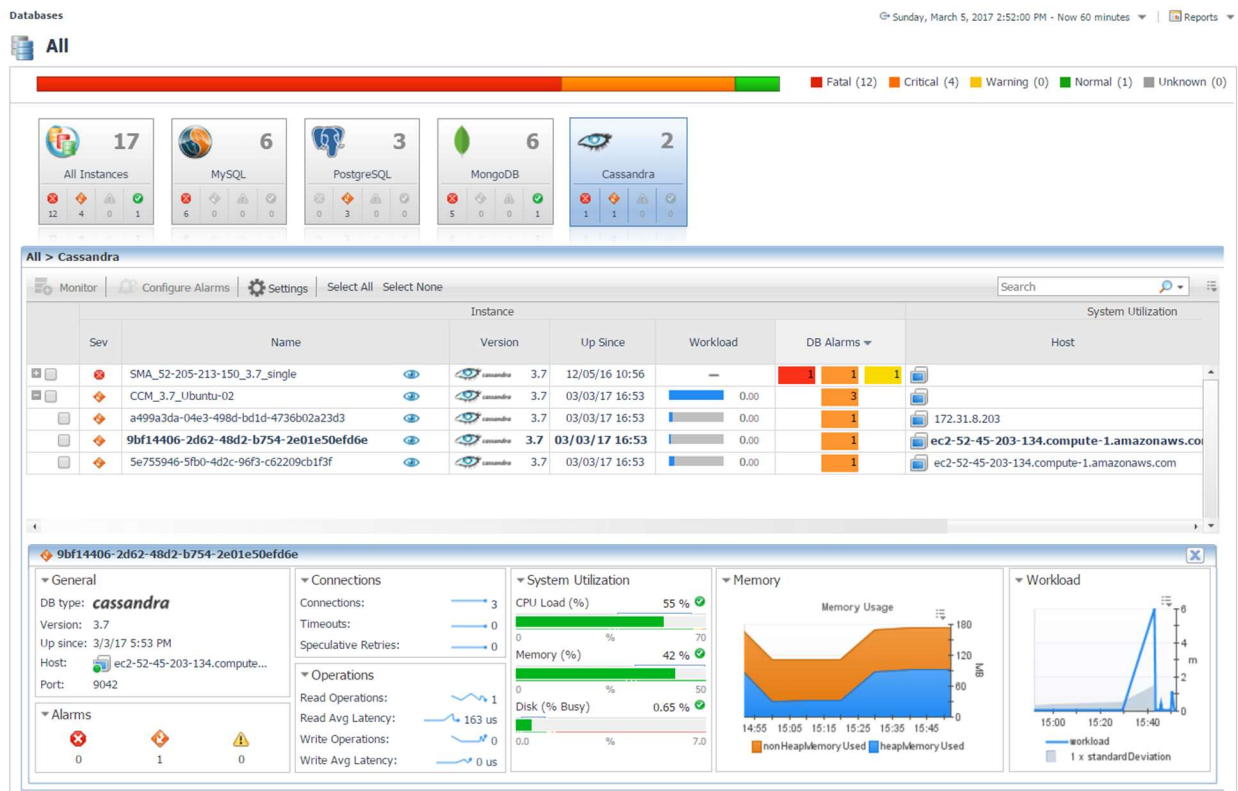
To create a registry variable:

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

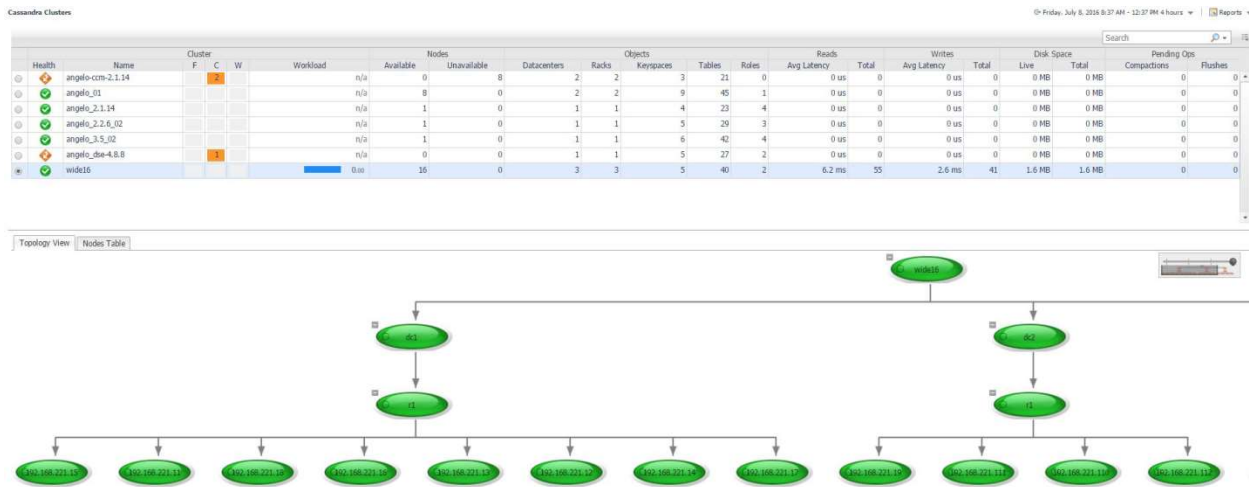
Databases

Foglight for Cassandra is now incorporated into the Databases dashboard along with any 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, alarms, and host utilization metrics. Clusters and Replica Sets can be expanded to show their individual nodes. Clicking on the eye icon in the Name column will bring up the Quick View with more key information. Clicking the name of an object will drill down to into the appropriate overview page.



Cassandra Clusters

This dashboard lists all monitored Cassandra Clusters and contains high-level information on cluster structure, nodes, health status, and key metrics. The workload metric is used for comparing the amount of work a node is doing. Cluster workload averages the workloads of all nodes in that cluster. Selecting a cluster will update the bottom section of the page and display either a topology view of the cluster structure or a list of nodes with relevant information. In the Topology View tab, hovering over an object will show a health summary, while clicking on a node will drill down to the Node Overview page, as will clicking on a node location in the Nodes Table. In the cluster table, clicking the Keyspace or Table column values will drill down to the Cluster Keyspaces and Cluster Tables pages, respectively. Other column values for metrics will show a time plot when hovered or clicked on.



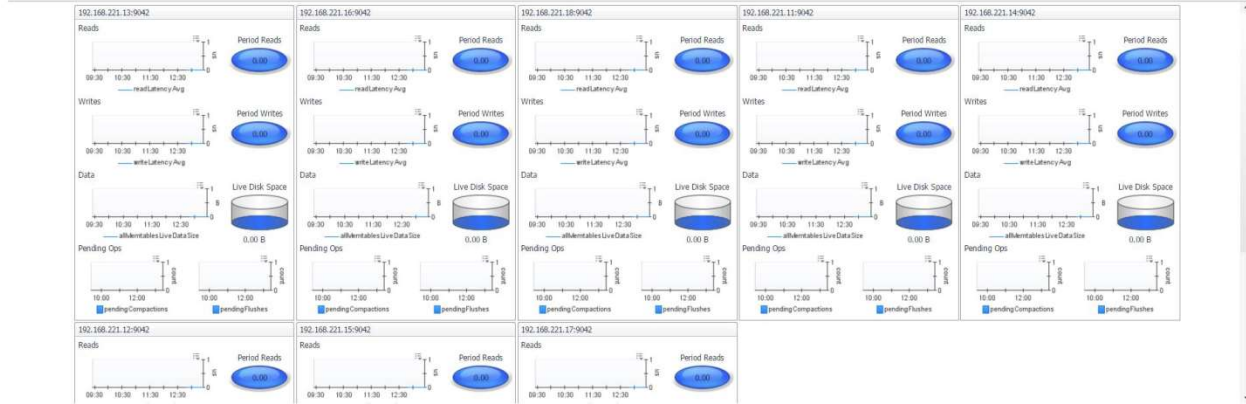
Cluster Keyspaces

This page lists all keyspaces in the selected cluster, aggregating metric information across nodes. Selecting a keyspace will update the section below, displaying a summary of that keyspace for each node in the cluster. Clicking the node location in the title of each summary will drill down to the Node Tables page for the selected node, filtering for only tables in that keyspace. In the Keyspace table, selecting the Tables column value will drill down to the Cluster Tables page, again filtering for only tables in that keyspace.

← Cassandra Clusters → ... → 192.168.221.11:9042 Keyspaces → widekeys Keyspaces

Friday, July 8, 2016 9:25 AM - 1:25 PM 4 hours Reports

| Health | Keyspace Name | F | C | W | Tables | Views | Reads | Writes | Disk Space | Pending Ops | Replication | Factor | Datcenters |
|--------|--------------------|---|---|---|--------|-------|-------------|--------|---------------|-------------|-------------|-------------|------------|
| | | | | | | | Avg Latency | Total | Live Memtable | Live | Total | Compactions | Flushes |
| OK | system | | | | 22 | 0 | 0/4.3 ms | 7403 | 2.2 MB | 1.9 MB | 1.9 MB | 0 | 0 |
| OK | system_auth | | | | 4 | 0 | 0/922.6 us | 7118 | 0 MB | 89.3 KB | 89.3 KB | 0 | 0 |
| OK | system_schema | | | | 30 | 0 | 0/0 us | 0 | 0 MB | 924.6 KB | 924.6 KB | 0 | 0 |
| OK | system_traces | | | | 2 | 0 | 0/0 us | 0 | 0 MB | 0 MB | 0 MB | 0 | 0 |
| OK | system_distributed | | | | 2 | 0 | 0/0 us | 0 | 0 MB | 0 MB | 0 MB | 0 | 0 |
| OK | widekeys | | | | 2 | 1 | 0/0 us | 0 | 0 MB | 0 MB | 0 MB | 0 | 0 |



Cluster Tables

This page lists tables in the selected cluster, aggregating metric information across nodes. Selecting a table will update the section below, displaying a summary of that table for each node in the cluster. Clicking the node location in the title of each summary will drill down to the Node Table page for the selected table. To filter the list of tables by keyspace, click the Select Keyspaces button at the top left of the table and select one or more keyspace for which you wish to view tables.



Traces Sample

The Traces page shows sampled entries from Cassandra's tracing system. Entries with the same query are aggregated and presented by average and maximum sampled duration. Individual query executions, called "sessions", are retrieved along with each execution duration, timestamp, source, consistency level, etc. The internal events generated to process the query are retrievable per session on demand. Tracing is not enabled by default, but nodetool can be used to enable tracing on a portion of all queries with its subcommand "settraceprobability".

Traces Sample - CCM_3.7_Ubuntu-02 Tuesday, October 17, 2017 1:34:05 PM - 2:34:05 PM 60 minutes Reports

<<Go to Clusters

| Sessions Sample | Request | Command | Query | Duration Avg | Duration Max | Sample Count |
|-----------------|--------------------|---------|--|--------------|--------------|--------------|
| View | Execute CQL3 query | QUERY | SELECT host_id FROM system.local; | 13.18 ms | 657.3 ms | 519 |
| View | Execute CQL3 query | QUERY | SELECT * FROM system.local WHERE key='local' | 12.52 ms | 334.81 ms | 211 |
| View | Execute CQL3 query | QUERY | SELECT key FROM system.local; | 11.15 ms | 194.98 ms | 210 |
| View | Execute CQL3 query | QUERY | SELECT keyspace_name,table_name,bloom_filter_fp_chance,caching,comment,compaction,compression... | 32.82 ms | 813.02 ms | 210 |
| View | Execute CQL3 query | QUERY | SELECT * FROM system_auth.roles; | 69.9 ms | 910.27 ms | 159 |
| View | Execute CQL3 query | QUERY | SELECT * FROM system_traces.sessions | 823.48 ms | 2.27 sec | 88 |
| View | Execute CQL3 query | QUERY | DELETE FROM tuna.skipjack WHERE sample=?; | 18.64 ms | 186.11 ms | 31 |
| View | Execute CQL3 query | QUERY | INSERT INTO salmon.sockeye (word,letter,time) VALUES (?,?); | 18.31 ms | 82.44 ms | 19 |
| View | Execute CQL3 query | QUERY | INSERT INTO tuna.bigeye (sample,country,river) VALUES (?,?); | 58.65 ms | 487.06 ms | 12 |
| View | Execute CQL3 query | QUERY | DELETE FROM tuna.albacore WHERE sample=?; | 16.44 ms | 60.15 ms | 9 |
| View | Execute CQL3 query | QUERY | INSERT INTO salmon.atlantic (sample,country,river) VALUES (?,?); | 14.67 ms | 35.53 ms | 9 |
| View | Execute CQL3 query | QUERY | DELETE FROM tuna.bigeye WHERE sample=?; | 38.25 ms | 134.67 ms | 9 |
| View | Execute CQL3 query | QUERY | INSERT INTO tuna.skipjack (sample,country,river) VALUES (?,?); | 59.36 ms | 167.34 ms | 7 |
| View | Execute CQL3 query | QUERY | INSERT INTO salmon.coho (sample,country,river) VALUES (?,?); | 22.92 ms | 65.45 ms | 7 |
| View | Execute CQL3 query | QUERY | DELETE FROM salmon.atlantic WHERE sample=?; | 47.32 ms | 205.16 ms | 7 |
| View | Execute CQL3 query | QUERY | INSERT INTO tuna.albacore (sample,country,river) VALUES (?,?); | 14.78 ms | 27.49 ms | 7 |
| View | Execute CQL3 query | QUERY | DELETE FROM salmon.chinook WHERE sample=?; | 10.62 ms | 29 ms | 6 |

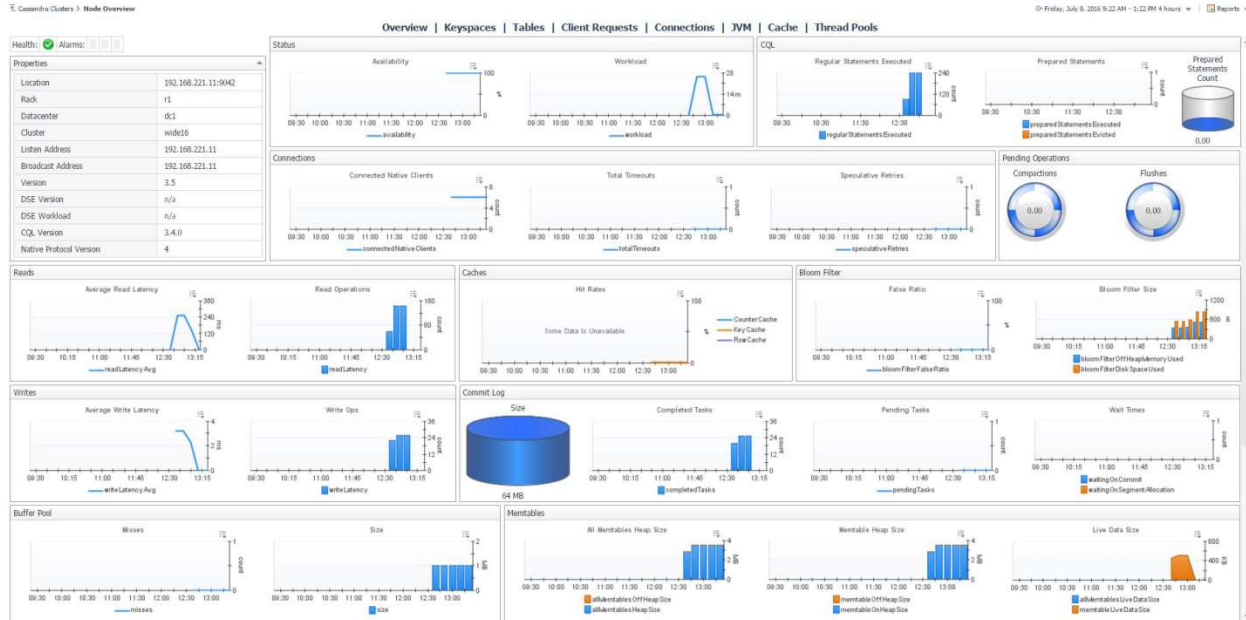
Traces Sample - CCM_3.7_Ubuntu-02 > Sessions Sample Tuesday, October 17, 2017 7:00 AM - 11:00 AM 4 hours Reports

| | | |
|---------|--------------------|---|
| Request | Execute CQL3 query | Query |
| Command | QUERY | SELECT keyspace_name,table_name,bloom_filter_fp_chance,caching,comment,compaction,compression,doclocal_read_repair_chance,default_time_to_live,gc_grace_seconds,memtable_flush_period_in_ms,read_repair_chance FROM system_schema.tables; |

| Sessions | | | | | | | Events | | |
|-------------------|------------|---------------|--------------|-----------|-------------|--------------------|--------|------------|--|
| Timestamp | Session ID | Client | Coordinator | Duration | Consistency | Serial Consistency | Load | Elapsed | Source |
| 10/17/17 10:18 AM | 1741565... | 172.31.52.138 | 172.31.8.201 | 819.15 ms | LOCAL_ONE | SERIAL | Load | 355 us | 172.31.8.201 Parsing SELECT keyspace_name,table_name,bloom_filter_fp_chi |
| 10/17/17 10:08 AM | b195ebf... | 172.31.52.138 | 172.31.8.201 | 582.55 ms | LOCAL_ONE | SERIAL | Load | 384 us | 172.31.8.201 Preparing statement |
| 10/17/17 10:48 AM | 3037e21... | 100.35.217.41 | 172.31.8.201 | 435.89 ms | LOCAL_ONE | SERIAL | Load | 442 us | 172.31.8.201 Computing ranges to query |
| 10/17/17 9:38 AM | 68da5f3... | 100.35.217.41 | 172.31.8.201 | 222.33 ms | LOCAL_ONE | SERIAL | Load | 1,255 us | 172.31.8.201 Submitting range requests on 1 ranges with a concurrency of 1 |
| 10/17/17 10:44 AM | a71b53e... | 172.31.52.138 | 172.31.8.201 | 182.32 ms | LOCAL_ONE | SERIAL | Load | 1,272 us | 172.31.8.201 Submitted 1 concurrent range requests |
| 10/17/17 9:14 AM | 2071b7f... | 100.35.217.41 | 172.31.8.201 | 177.13 ms | LOCAL_ONE | SERIAL | Load | 1,300 us | 172.31.8.201 Executing seq scan across 3 sstables for (min(-92233720368547 |
| 10/17/17 10:25 AM | 0b8ed75... | 100.35.217.41 | 172.31.8.201 | 172.67 ms | LOCAL_ONE | SERIAL | Load | 181,660 us | 172.31.8.201 Read 41 live and 0 tombstone cells |
| 10/17/17 9:25 AM | a1bb956... | 172.31.52.138 | 172.31.8.201 | 150.07 ms | LOCAL_ONE | SERIAL | Load | | |
| 10/17/17 9:54 AM | a50ee64... | 100.35.217.41 | 172.31.8.201 | 139.43 ms | LOCAL_ONE | SERIAL | Load | | |
| 10/17/17 9:16 AM | 6772ecc... | 100.35.217.41 | 172.31.8.201 | 117.37 ms | LOCAL_ONE | SERIAL | Load | | |
| 10/17/17 10:38 AM | dc76cd4... | 100.35.217.41 | 172.31.8.201 | 109.71 ms | LOCAL_ONE | SERIAL | Load | | |

Node Overview

This page provides a comprehensive view of the Cassandra Node, with health and alarms and configuration information at the top left and metrics grouped into relevant categories on the rest of the page. Once in the node section of the dashboards, the navigation bar at the top of the page can be used to navigate between different pages containing more information on the selected node. The Node Selector in the action panel on the right can also be used to switch between nodes in the same cluster.



Node Keyspaces

This page lists all keyspaces on the selected node, with metric information on reads and writes, latency timing, disk space, and pending operations. Clicking the Keyspace Name or Tables column value will drill down to the Node Tables page, filtering for only tables in that keyspace. The “Go to Cluster Keyspaces” link above the table will link to that page for the same cluster, where you can compare keyspaces across different nodes by selecting a keyspace row.

Overview | Keyspaces | Tables | Client Requests | Connections | JVM | Cache | Thread Pools

<<< Go to Cluster Keyspaces

| Health | Keyspace Name | F | C | W | Tables | Reads | | Writes | | Cassandra Latencies | | | Live Memtable | Disk Space | | Pending Ops | | |
|--------|--------------------|---|---|---|--------|-------------|---------|-------------|--------|---------------------|---------|--------|---------------|------------|----------|-------------|---------|---|
| | | | | | | Avg Latency | Total | Avg Latency | Total | Prepare | Propose | Commit | | Live | Total | Compactions | Flushes | |
| ✓ | system_traces | | | | 2 | | 0.0 us | 0 | 0.0 us | 0 | 0.0 us | 0.0 us | 0.0 us | 0 B | 0 B | 0 B | 0 | 0 |
| ✓ | system | | | | 22 | | 13.3 ms | 1 | 0.0 us | 0 | 0.0 us | 0.0 us | 0.0 us | 279 KB | 275.5 KB | 275.5 KB | 0 | 0 |
| ✓ | system_distributed | | | | 2 | | 0.0 us | 0 | 0.0 us | 0 | 0.0 us | 0.0 us | 0.0 us | 0 B | 0 B | 0 B | 0 | 0 |
| ✓ | system_schema | | | | 10 | | 0.0 us | 0 | 0.0 us | 0 | 0.0 us | 0.0 us | 0.0 us | 0 B | 133 KB | 133 KB | 0 | 0 |
| ✓ | system_auth | | | | 4 | | 0.0 us | 0 | 0.0 us | 0 | 0.0 us | 0.0 us | 0.0 us | 0 B | 4.9 KB | 4.9 KB | 0 | 0 |
| ✓ | widekeys | | | | 2 | | 0.0 us | 0 | 0.0 us | 0 | 0.0 us | 0.0 us | 0.0 us | 0 B | 0 B | 0 B | 0 | 0 |

Node Tables

This page lists tables on the selected node, aggregating metric information across nodes, with metric information on reads and writes, latency timing, disk space, and pending operations. Clicking the Table Name column value will drill down to the Node Table page for more information on that table. The “Go to Cluster Tables” link above the table will link to that page for the same cluster, where you can compare tables across different nodes by selecting a table row. To filter the list of tables by keyspace, click the Select Keyspaces button at the top left of the table and select one or more keyspace for which you wish to view tables.

Overview | Keyspaces | Tables | Client Requests | Connections | JVM | Cache | Thread Pools

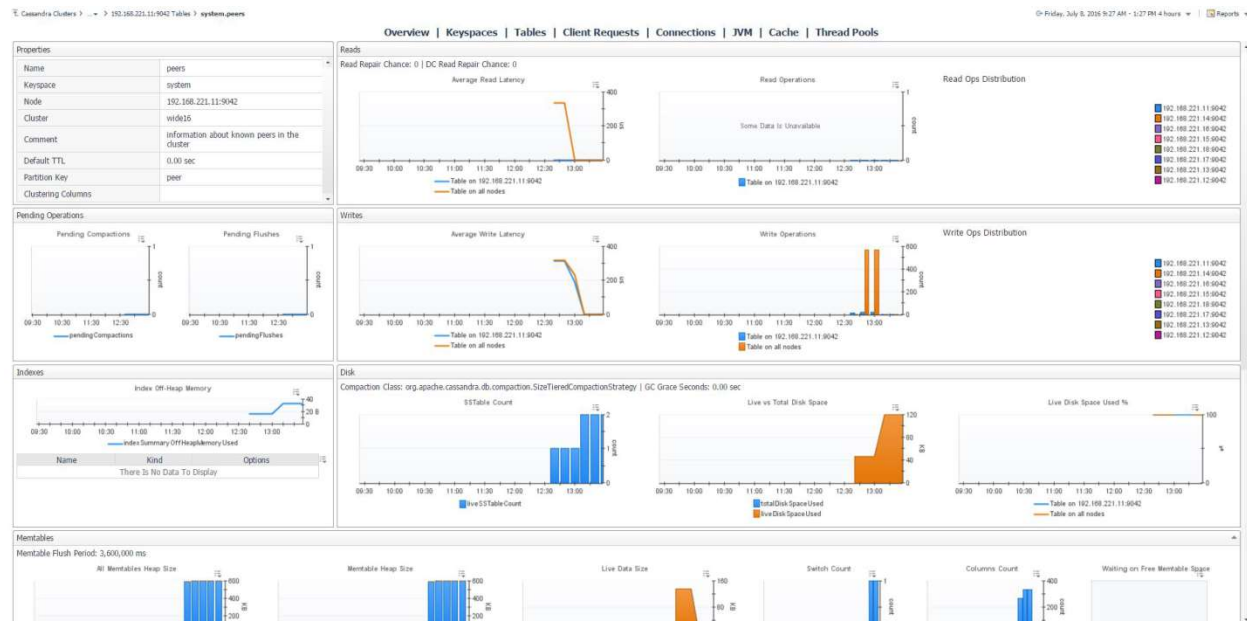
<<< Go to Cluster Tables

Select Keyspaces

| Health | Name | Table | Keyspace | F | C | W | Objects | Views | Avg Latency | Reads | Total | Avg Latency | Writes | Total | Key Cache | Row Cache | Live Memtable | Disk Space | Live | Total | Pending Ops | Comps | Flushes |
|--------|-----------------------|-------|----------|---|---|---|---------|-------|-------------|-------|-------|-------------|--------|-------|-----------|-----------|---------------|------------|----------|-------|-------------|-------|---------|
| ✓ | local | | system | | | | 0 | 0 | 11.4 ms | 1 | 1 | 0.0 us | 0 | 0 | 1.0 % | 0.0 % | 0 B | 0 B | 25 KB | 0 | 0 | | |
| ✓ | schema_columnfamilies | | system | | | | 0 | 0 | 0.0 us | n/a | n/a | 0.0 us | n/a | n/a | n/a | n/a | n/a | n/a | n/a | n/a | n/a | n/a | n/a |
| ✓ | schema_usertypes | | system | | | | 0 | 0 | 0.0 us | n/a | n/a | 0.0 us | n/a | n/a | n/a | n/a | n/a | n/a | n/a | n/a | n/a | n/a | n/a |
| ✓ | peers | | system | | | | 0 | 0 | 0.0 us | 0 | 0 | 0.0 us | 0 | 0 | 0.0 % | 0.0 % | 0 B | 0 B | 119.4 KB | 0 | 0 | | |
| ✓ | size_estimates | | system | | | | 0 | 0 | 0.0 us | 0 | 0 | 0.0 us | 0 | 0 | 0.0 % | 0.0 % | 278.9 KB | 278.9 KB | 87.8 KB | 0 | 0 | | |
| ✓ | range_refers | | system | | | | 0 | 0 | 0.0 us | 0 | 0 | 0.0 us | 0 | 0 | 0.0 % | 0.0 % | 0 B | 0 B | 0 B | 0 | 0 | | |
| ✓ | compaction_history | | system | | | | 0 | 0 | 0.0 us | 0 | 0 | 0.0 us | 0 | 0 | 0.0 % | 0.0 % | 0 B | 0 B | 13 KB | 0 | 0 | | |

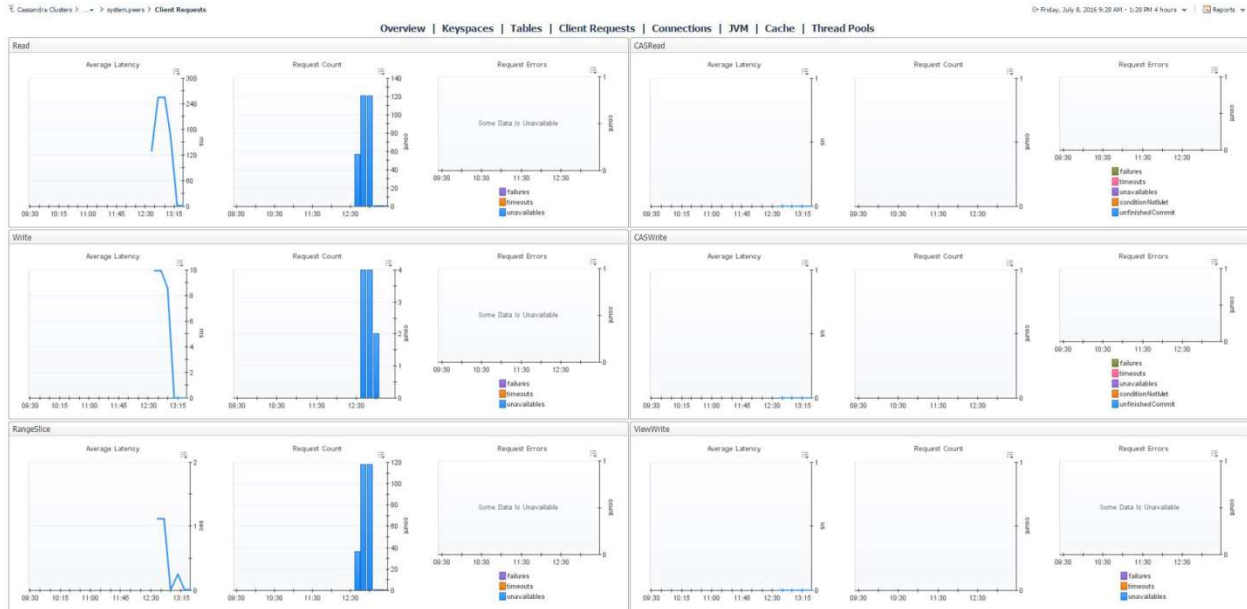
Node Table

This page provides a comprehensive view of the table, with configuration information at the top left and metrics grouped into relevant categories on the rest of the page.



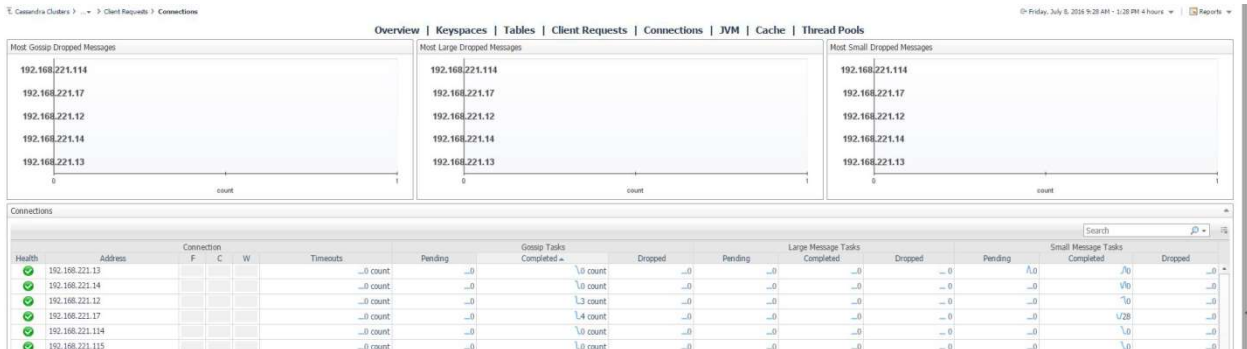
Client Requests

This page displays client requests to the node by type, showing metric histories for average latency, request counts, and request errors, broken down by error type.



Connections

The Connections page shows connections between the selected node and other nodes in the cluster. At the top of the page, three bar graphs show nodes with the most gossip, large, and dropped messages. The below table lists all connections, showing timeouts and result status of messages by task type.



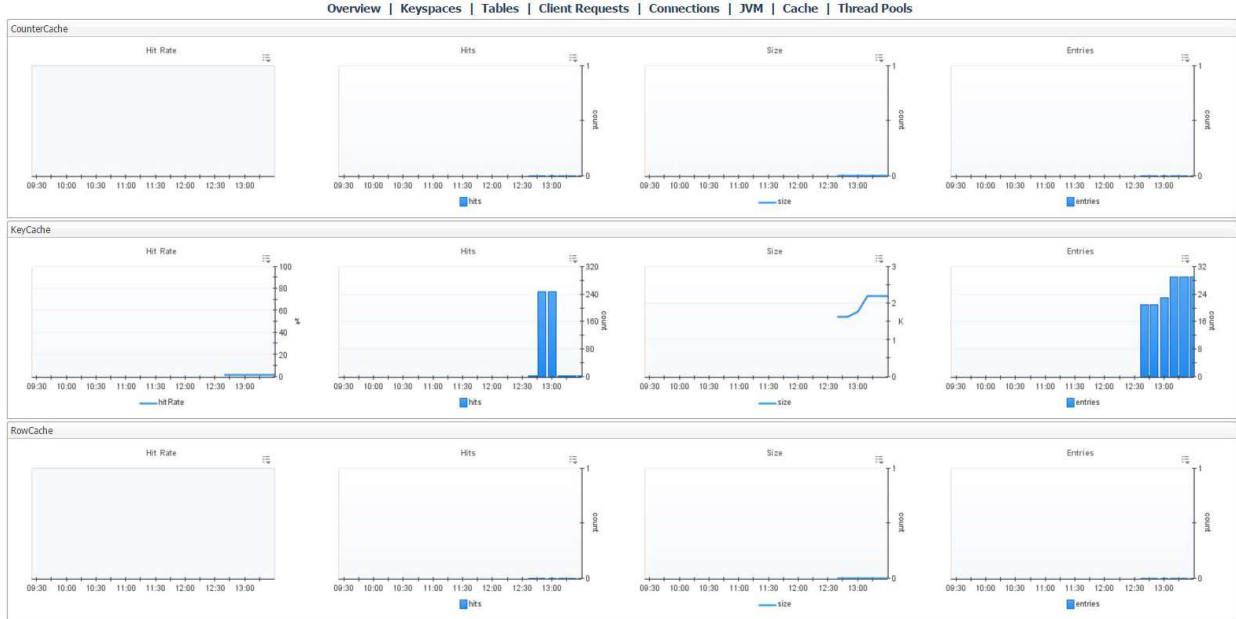
JVM

This page features information on the JVM that Cassandra runs on. The top row contains JVM properties, uptime, and information on memory and threading. The second row displays OS information and memory and CPU metrics. The final row features operational metrics for the various buffer pools, garbage collectors, and memory pools maintained by the JVM.



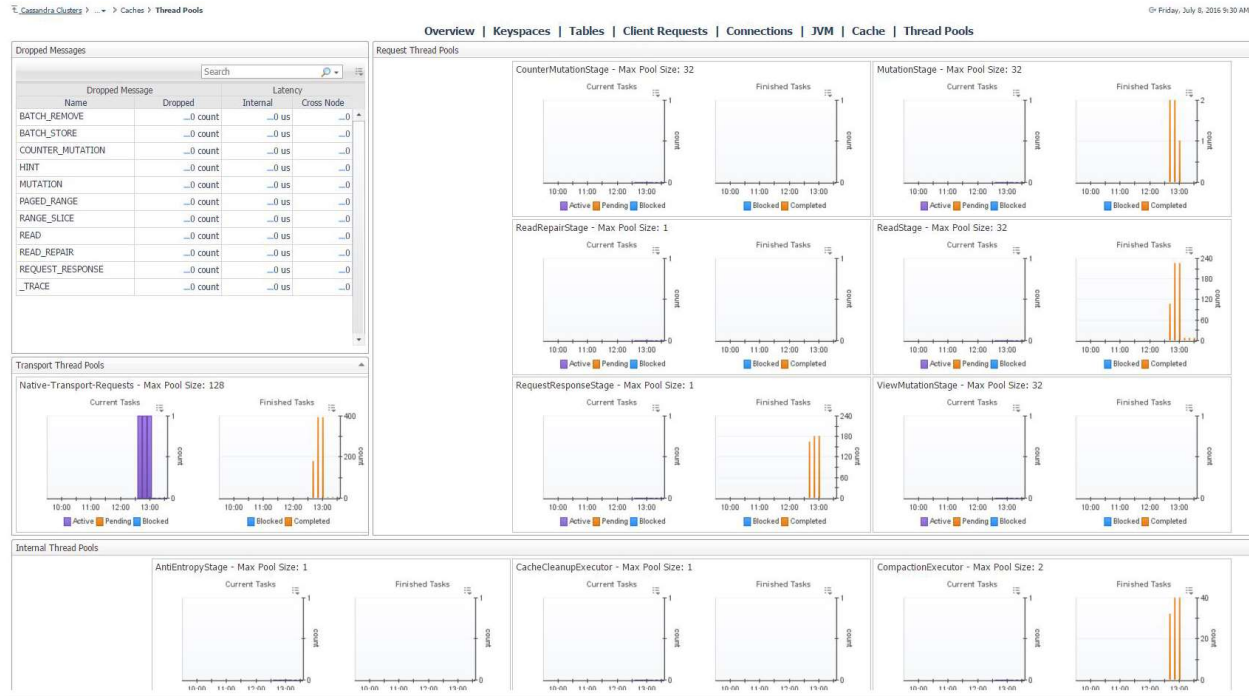
Cache

The Cache page features cache-related metrics including hit rate, hits, size and number of entries for the counter cache, key cache, and row cache.



Thread Pools

This page displays active and completed task statuses for every thread pool type in the Cassandra node, grouped into Request, Transport, and Internal categories. The page also features a table of dropped messages by stage and includes metrics for number of dropped messages and latencies for internal and cross node dropped messages.



Slow Query Log

The slow query log dashboard displays information on long running queries for DSE Cassandra servers. Slow query logging is enabled by default. The slow query threshold can be raised or lowered with the dsetool "perf cqlslowlog" subcommand.

Slow Query Log Table:

| Timestamp | Source IP | Username | Duration | Tables | Command(s) |
|------------------|--------------|----------|----------|-----------------|---|
| 10/17/17 9:17 AM | 172.31.8.201 | root | 20 ms | slowspace.place | SELECT * FROM slowspace.place; |
| 10/17/17 9:17 AM | 172.31.8.201 | root | 2 ms | slowspace.place | SELECT * FROM slowspace.place; |
| 10/17/17 9:17 AM | 172.31.8.201 | root | 20 ms | slowspace.place | SELECT * FROM slowspace.place; |
| 10/17/17 9:17 AM | 172.31.8.201 | root | 2 ms | slowspace.place | INSERT INTO slowspace.place(id, ts, info, note) VALUES(uuid(), dateof(now()), 'some info', 'a note'); |
| 10/17/17 9:17 AM | 172.31.8.201 | root | 31 ms | slowspace.place | INSERT INTO slowspace.place(id, ts, info, note) VALUES(uuid(), dateof(now()), 'some info', 'a note'); |
| 10/17/17 9:17 AM | 172.31.8.201 | root | 86 ms | slowspace.place | INSERT INTO slowspace.place(id, ts, info, note) VALUES(uuid(), dateof(now()), 'some info', 'a note'); |

Rules

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 unavailables for client request reads and writes is high relative to recent averages.

Cassandra CommitLog PendingTasks

Alert if the number of commit log pending tasks is too high.

Cassandra CommitLog WaitingOn

Alert if the number of commit log waiting on segment allocation and/or waiting on commit is too high.

Cassandra Compaction PendingTasks

Alert if the number of compaction pending tasks is too high.

Cassandra Datacenter Availability

Alert if the percent of available nodes in a datacenter is too low.

Cassandra DroppedMessage MutationDropped

Alert if there were any dropped mutation messages.

Cassandra JVM Memory HeapMemoryRatio

Alert if the JVM heap memory ratio of used memory to max memory is high.

Cassandra JVM OS SystemCPULoad

Alert if the JVM operating system CPU load is too high.

Cassandra JVM OS UsedPhysicalMemSizePct

Alert if the JVM operating system used physical memory size percent is too high.

Cassandra JVM OS UsedSwapSpaceSizePct

Alert if the JVM operating system used swap space size percent is too high.

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 too high.

Cassandra Node LiveDiskSpaceUsedPercent

Alert if the used percent of live disk space on a node is too low.

Cassandra Node PendingCompactions

Alert if the number of pending compactions on a node is too high.

Cassandra Node PendingFlushes

Alert if the number of pending flushes on a node is too high.

Cassandra Node RowCacheHitRate

Alert if the row cache hit rate on a node is too low.

Cassandra SchemaVersion

Alert if schema version is not identical for all nodes.

Cassandra Storage Exceptions

Alert if there are storage exceptions on a node.

Cassandra Storage LoadVsSpace

Alert if the current storage load is large relative to available disc space on a node. Only available in conjunction with the Infrastructure cartridge.

Cassandra Storage TotalHints

Alert if the number of storage total hints on a node is too high.

Cassandra Storage TotalHintsInProgress

Alert if the number of storage total hints in progress on a node is too high.

Cassandra Table AvgKeyCacheHitRate

Alert if the average key cache hit rate for a table across all nodes is low.

Cassandra Table AvgReadLatency

Alert if the average read latency on a table is too high.

Cassandra Table AvgRowCacheHitRate

Alert if the average row cache hit rate for a table across all nodes is low.

Cassandra ThreadPools PendingTasks

Alert if the number of thread pool pending tasks is high.

Cassandra ThreadPools TotalBlockedTasks

Alert if the number of thread pool total blocked tasks is high.

Reports

Cassandra Cluster Summary

Summary of a Cassandra Cluster with availability, disk space, operations, nodes, and largest keyspaces.

Cassandra Cluster Tables

Top Cassandra Cluster Tables by reads, writes, avg read/write latency, or disk space.

Cassandra Node Executive Summary

Executive summary of a Cassandra node with availability and connection info, workload and JVM resource usage, and top alarms.

Cassandra Node Health Check

Health check report for a Cassandra node with availability and connection info, workload and JVM resource usage, and top statements, alarms, and tables.

Cassandra Storage Report

Shows Cassandra server storage capacity, growth rate, etc. Note: Host monitoring must be enabled in order to retrieve space remaining and days remaining until full. Report on all servers that are projected to fill up before this many days. Report on servers with less than this percentage of remaining disk space.