



Red Hat Gluster Storage 3.5

Monitoring Guide

Monitoring Gluster Cluster

Red Hat Gluster Storage 3.5 Monitoring Guide

Monitoring Gluster Cluster

Legal Notice

Copyright © 2021 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This guide provides essential information on how to import a Gluster cluster using Red Hat Gluster Storage Web Administration and how to monitor your Gluster cluster health, performance, and status. The monitoring and metrics visualization is provided by the Grafana monitoring platform which is integrated into the Web Administration interface. Making open source more inclusive Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see our CTO Chris Wright's message

Table of Contents

CHAPTER 1. OVERVIEW	4
CHAPTER 2. IMPORT CLUSTER	5
2.1. IMPORTING CLUSTER	5
2.1.1. Troubleshooting Import Cluster	8
2.1.2. Volume Profiling	8
CHAPTER 3. CLUSTER EXPANSION	12
3.1. TROUBLESHOOTING CLUSTER EXPANSION	14
CHAPTER 4. UNMANAGING CLUSTER	15
4.1. GRAPHITE AND CARBON DATA ARCHIVAL	16
4.2. GRAPHITE DATABASE METRICS RETENTION	16
CHAPTER 5. WEB ADMINISTRATION LOG LEVELS	18
CHAPTER 6. MONITORING AND METRICS	20
CHAPTER 7. MONITORING DASHBOARD AND CONCEPTS	21
7.1. DASHBOARD SELECTOR	21
7.2. DASHBOARD PANELS	21
7.3. DASHBOARD ROWS	23
7.4. DASHBOARD COLOR CODES	24
CHAPTER 8. MONITORING DASHBOARD FEATURES	25
8.1. DASHBOARD SEARCH	25
8.2. DASHBOARD TIME RANGE	25
8.3. DASHBOARD SHARING	26
CHAPTER 9. MONITORING DASHBOARD NAVIGATION	27
9.1. CLUSTER VIEW DASHBOARD	27
9.2. HOSTS VIEW DASHBOARD	27
9.3. VOLUME VIEW DASHBOARD	28
9.4. BRICK VIEW DASHBOARD	28
CHAPTER 10. MONITORING CLUSTER METRICS	29
10.1. CLUSTER LEVEL DASHBOARD	29
10.1.1. Monitoring and Viewing Cluster Health	29
10.1.1.1. Health and Snapshots	29
10.1.1.2. Hosts, Volumes and Bricks	29
10.1.1.3. Geo-Replication Session	30
10.1.1.4. Health Panel Descriptions	31
10.1.2. Monitoring and Viewing Cluster Performance	31
10.1.3. Top Consumers	34
10.1.4. Monitoring and Viewing Cluster Status	37
10.2. HOST LEVEL DASHBOARD	39
10.2.1. Monitoring and Viewing Health and Status	39
10.2.2. Monitoring and Viewing Performance	40
10.2.2.1. Memory and CPU Utilization	40
10.2.2.2. Capacity and Disk Load	43
10.2.2.3. Network	47
10.2.3. Host Dashboard Metric Units	49
10.3. VOLUME LEVEL DASHBOARD	50
10.3.1. Monitoring and Viewing Health	50

10.3.2. Monitoring and Viewing Performance	54
10.3.3. Monitoring File Operations	59
10.3.4. Volume Dashboard Metric Units	61
10.4. BRICK LEVEL DASHBOARD	62
10.4.1. Monitoring and Viewing Brick Status	62
10.4.2. Monitoring and Viewing Brick Performance	63
10.4.3. Brick Dashboard Metric Units	68
CHAPTER 11. USERS AND ROLES ADMINISTRATION	70
11.1. USER ROLES	70
11.2. CONFIGURING ROLES	70
11.2.1. Editing Users	71
11.2.2. Disabling Notifications and Deleting User	72
CHAPTER 12. ALERTS AND NOTIFICATIONS	74
12.1. TYPES OF ALERTS	74
12.2. LIST OF ALERTS	74
12.3. ALERTS NOTIFICATIONS DRAWER	75
12.4. SMTP NOTIFICATIONS CONFIGURATION	76
12.5. SNMPV3 NOTIFICATION CONFIGURATION	77

CHAPTER 1. OVERVIEW

Red Hat Gluster Storage Web Administration provides visual monitoring and metrics infrastructure for Red Hat Gluster Storage 3.5 and is the primary method to monitor your Red Hat Gluster Storage environment. The Red Hat Gluster Storage Web Administration is based on the Tendrl upstream project and utilizes Ansible automation for installation. The key goal of Red Hat Gluster Storage Web Administration is to provide deep metrics and visualization of Red Hat Storage Gluster clusters and the associated physical storage elements such as storage nodes, volumes, and bricks.

Key Features

1. Monitoring dashboards for Clusters, Hosts, Volumes, and Bricks
2. Top-level list views of Clusters, Hosts, and Volumes
3. SNMPv3 Configuration and alerting
4. User Management
5. Importing Gluster cluster

CHAPTER 2. IMPORT CLUSTER

You can import and start managing a cluster once Web Administration is installed and ready. For installation instructions, see the [Installing Web Administration](#) chapter of the Red Hat Gluster Storage Web Administration Quick Start Guide.

2.1. IMPORTING CLUSTER

The following procedure outlines the steps to import a Gluster cluster.

Procedure. Importing Cluster

1. Log in to the Web Administration interface. Enter the username and password, and click **Log In**.



NOTE

The default username is **admin** and the default password is **adminuser**.

Figure 2.1. Login Page



- In the default landing interface, a list of all the clusters ready to be imported is displayed. Locate the cluster to be imported and click **Import**.

Figure 2.2. Import Cluster

Clusters

Name ▾ *Filter by Name* | Name ▾ *↓*

1 Cluster

	Cluster Version	Managed	Hosts	
? 0354fa28-78b4-47c4-a960-4aa1130d8771	RHGS 3.4.0	No	6	Import

- Enter a user-friendly Cluster name. By default, the **Enable for all volumes** option is selected.

Figure 2.3. Cluster name

[Clusters](#) » Import Cluster

Import Cluster

Cluster Name

Only alphanumeric and underscore characters are allowed.

Ensure all hosts in the cluster are listed below, then click **Import** to begin importing the cluster.

Volume Profiling

- Enable for all volumes
- Disable for all volumes
- Keep current configuration

Discovered Host(s)

Name ▾ *Filter by Name*

6 of 6 Discovered - Hosts -



NOTE

If the cluster name is not provided, the system will assign a randomly generated uuid as a cluster name. However, it is advisable to enter a user-friendly cluster name to easily locate the cluster from the clusters list.

- Click **Import** to continue.

6 of 6 Discovered - Hosts -

Host ^
tendrl-usm1-gl1.example.org
tendrl-usm1-gl2.example.org
tendrl-usm1-gl3.example.org
tendrl-usm1-gl4.example.org
tendrl-usm1-gl5.example.org
tendrl-usm1-gl6.example.org



1. The cluster import request is submitted. To view the task progress, click **View Task Progress**.

Figure 2.4. Task Detail



Import Cluster Task Submitted

An import cluster task has been submitted. You will be notified when processing is complete and the cluster is ready for use.




2. Navigate to the **All Clusters** interface view. The Cluster is successfully imported and ready for use.

Figure 2.5. Cluster Ready

Clusters

Name ▾ | Name ▾ ↓↑

1 Cluster

	Cluster Version	Managed	Hosts	Volumes	Alerts	Volume Profiling	
 ClusterA	RHGS 3.4.0	Yes	6	3	1	Enabled	Ready to Use

2.1.1. Troubleshooting Import Cluster

Scenario: The Import cluster UI button is disabled after a failed cluster import operation.

In this scenario, when cluster import fails, the Import button is disabled.

Resolution

Resolve the issue by investigating why the import cluster operation failed. To see details of the failed operation, navigate to the **All Clusters** interface of the Web Administration environment. In the clusters list, locate the cluster that you attempted to import and click on **View Details** next to the **Import Failed** status label.

Examine the reason of the failed cluster import operation and resolve the issue. After resolving the issue, unmanage the cluster and then reimport the cluster. For unmanaging cluster instructions, navigate to the [Unmanaging Cluster](#) section of this Guide.

2.1.2. Volume Profiling

Volume profiling enables additional telemetry information to be collected on a per volume basis for a given cluster, which helps in troubleshooting, capacity planning, and performance tuning.

Volume profiling can be enabled or disabled on a per cluster basis and per volume basis when a cluster is actively managed and monitored using the Web Administration interface.




NOTE

Enabling volume profiling results in richer set of metrics being collected which may cause performance degradation to occur as system resources, for example, CPU and memory, may get used for volume profiling data collection.

Volume Profiling at Cluster Level

To enable or disable volume profiling at cluster level:

1. Log in to the Web Administration interface
2. From the Clusters list, locate the cluster to disable Volume Profiling.

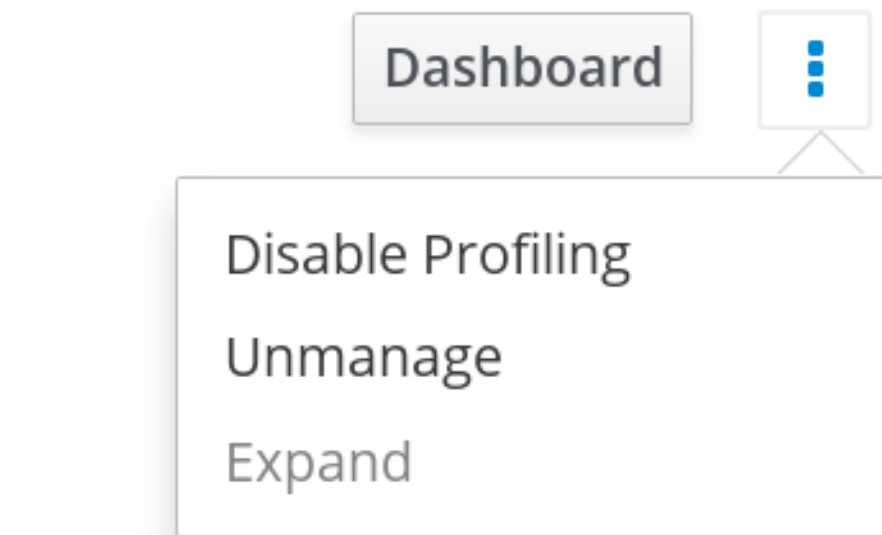


NOTE

Clusters list is the default landing interface after login and the Interface switcher is on **All Clusters**.

3. At the right-hand side, next to the Dashboard button, click the vertical ellipsis.
4. An inline menu is opened. Click **Disable Profiling** or **Enable Profiling** depending on the current state. In the example screen below, Volume Profiling option is enabled. Click **Disable Profiling** to disable.

Figure 2.6. Disable Volume Profiling



5. The disable profiling task is submitted and processed. After processing, Volume Profiling is successfully disabled.

Figure 2.7. Disable Volume Profiling

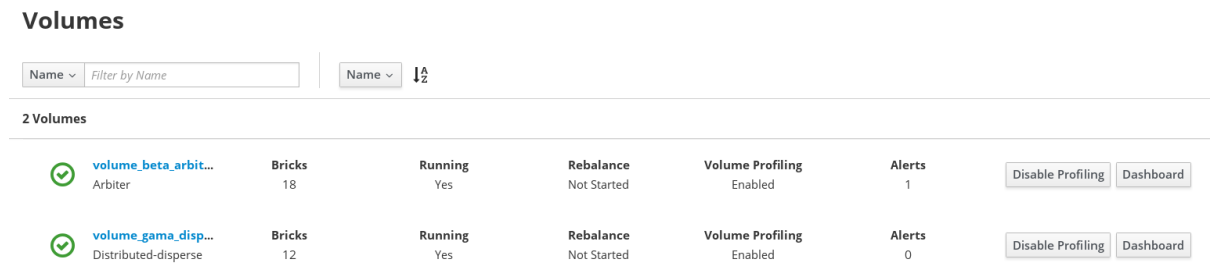
Hosts	Volumes	Alerts	Volume Profiling
4	1	0	Disabled

Volume Profiling at Volume Level

To enable or disable volume profiling at Volume level:

1. Log in to the Web Administration interface and select the specific cluster from the Interface switcher drop-down.
2. After selecting the specific cluster, the left vertical navigation pane is exposed.
3. From the navigation pane, click **Volumes**. The Volumes view is displayed listing all the Volumes part of the cluster.

Figure 2.8. Volumes View




- Locate the volume and click **Disable Profiling** or **Enable Profiling** depending on the current state. In the example screen below, Volume Profiling is enabled. To disable volume profiling, click **Disable Profiling**.

Figure 2.9. Disable Volume Profiling



- The disable profiling task is submitted and processed. After processing, Volume Profiling is successfully disabled.

Figure 2.10. Disable Volume Profiling



Volume Profiling Metrics

When volume profiling is disabled, the following metrics will not be displayed in the Grafana Dashboard. Based on the metrics required to view, enable or disable volume profiling accordingly.

For detailed information on Volume Profiling, see the [Monitoring Red Hat Gluster Storage Gluster Workload](#) chapter of the Red Hat Gluster Storage Administration Guide.

Table 2.1. Volume Profiling Metrics

Grafana Dashboard Level	Dashboard Section	Panel and Metrics
Cluster Dashboard	At-a-glance	IOPS
Host Dashboard	At-a-glance	Brick IOPS
Volume Dashboard	Performance	IOPS
Volume Dashboard	Profiling Information	File Operations For Locks
Volume Dashboard	Profiling Information	Top File Operations

Grafana Dashboard Level	Dashboard Section	Panel and Metrics
Volume Dashboard	Profiling Information	File Operations for Read/Write
Volume Dashboard	Profiling Information	File Operations for Inode Operations
Volume Dashboard	Profiling Information	File Operations for Entry Operations
Brick Dashboard	At-a-glance	IOPS

CHAPTER 3. CLUSTER EXPANSION

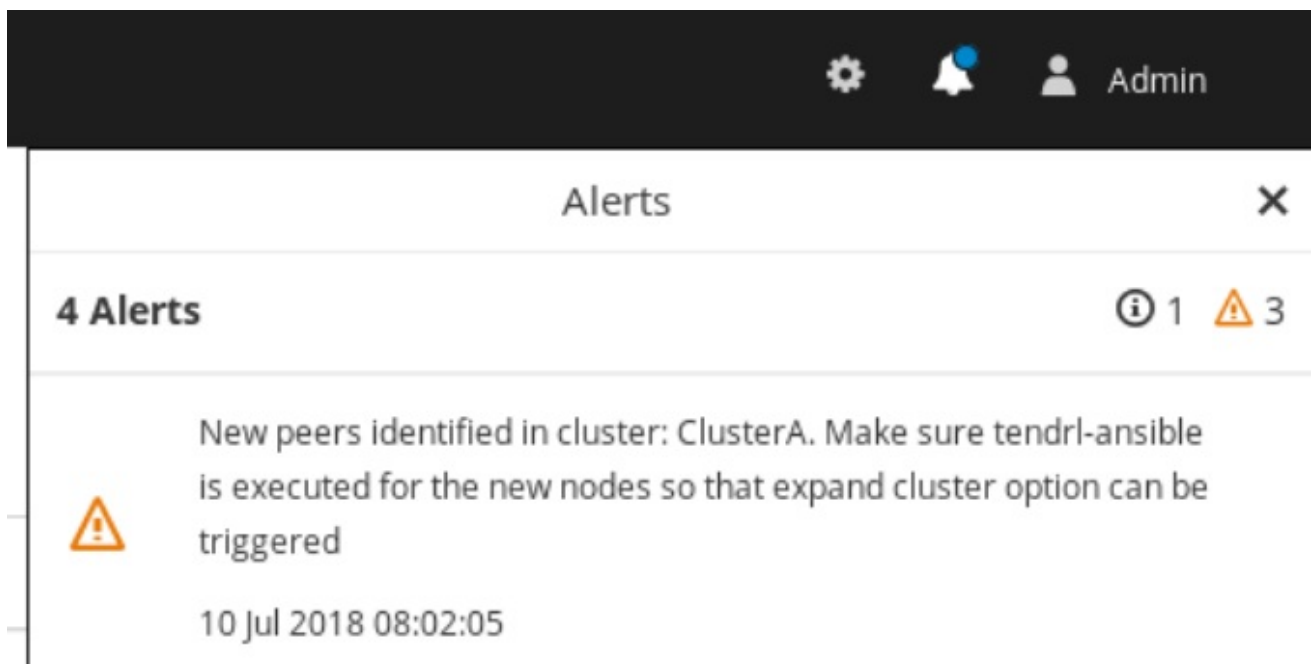
To expand an existing Gluster cluster already imported and managed by the Web Administration environment, perform the following sequence of actions:

Create New Nodes

Before initiating cluster expansion in the Web Administration interface, create new Gluster storage nodes to be imported in Web Administration. For detailed instructions, see the [Expanding Volumes](#) section in the *Red Hat Gluster Storage Administration Guide*.

New Alert

After successfully expanding the cluster with new nodes, an alert is generated in the Notifications panel of the Web Administration interface as displayed in the following screen:



In this screen, new peers represents new storage nodes that were created in the previous cluster expansion process.

New Event

Additionally, a new Event is also generated about the availability of new storage nodes for expansion as displayed in the following screen:



Run Tendrl-ansible installation for new nodes

After the storage nodes are identified by Web Administration, run `tendr-ansible` installation by adding the new hosts in the inventory file. For detailed installation instructions, see the [Installing Web Administration](#) chapter in the Red Hat Gluster Storage Web Administration Quick Start Guide.



IMPORTANT

If you have implemented TLS-based client-server authentication, install and deploy TLS encryption certificates on the new nodes before running `tendr-ansible` installation to avoid cluster expansion failure. For detailed information on configuring TLS encryption, see the [TLS Encryption Configuration](#) chapter of the Red Hat Gluster Storage Web Administration Quick Start Guide.

Expand Cluster in Web Administration Environment

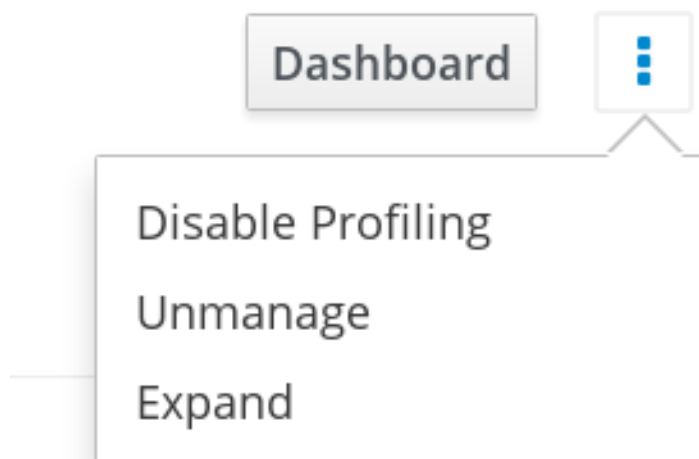
After `tendr-ansible` is executed, the expand cluster option becomes available in the Web Administration interface. Follow these instructions to expand the cluster:

Procedure. Expanding Cluster

1. Log in to the Web Administration interface and click **All Clusters** interface from the interface selector drop-down.
2. In the Clusters view, locate the cluster to be expanded. It will be labelled as **Expansion required**.

Alerts	Volume Profiling	Expansion required
3	Enabled	

3. At the far right of the cluster row, click the vertical ellipsis and then click **Expand**.



4. The available hosts to be added are listed. Click **Expand**.

Expand Cluster
✕

ClusterA will be expanded by importing the new hosts listed below. Further actions will not be allowed on this cluster until expansion is complete.


Name ▾

2 Hosts Discovered

Host ^	Address
tendrl-usm1-gl5.example.org	10.37.169.27
tendrl-usm1-gl6.example.org	10.37.169.224

Cancel
Expand

5. The cluster expansion task is submitted.



6. The cluster is successfully expanded and ready for use.

Cluster Version	Managed	Hosts	Volumes	Alerts	Volume Profiling	Ready to Use
RHGS 3.5	Yes	6	2	2	Enabled	

3.1. TROUBLESHOOTING CLUSTER EXPANSION

Scenario: Cluster Expansion task failed

In this scenario, the cluster expansion task fails.

Resolution

If cluster expansion fails, check if tendrl-ansible was executed successfully and ensure the node agents are correctly configured. If cluster expansion failed due to errors, resolve the errors on affected nodes and re-initiate the Expand Cluster action.

To verify tendrl-ansible execution steps, see the [Web Administration Installation](#) section of the Red Hat Gluster Storage Web Administration Quick Start Guide.

For detailed instructions on expanding cluster, navigate to the [Cluster Expansion](#) section of this Guide.

CHAPTER 4. UNMANAGING CLUSTER

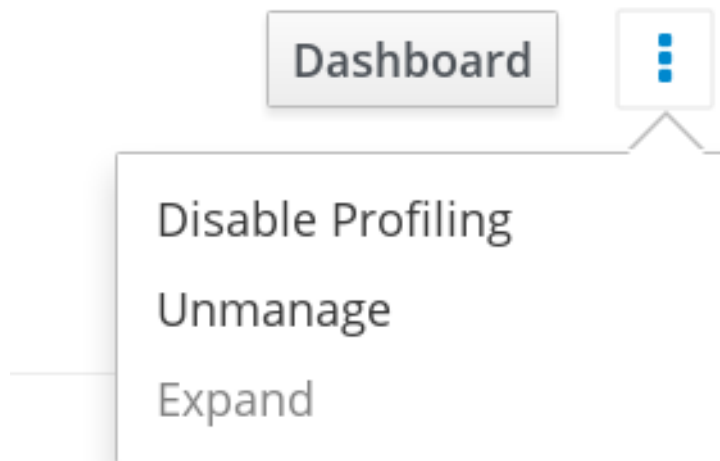
The following are the reasons to unmanage a cluster:

1. Import cluster fails
2. Removal of a cluster that is no longer under management and used for monitoring
3. A way to remove orphaned cluster entries

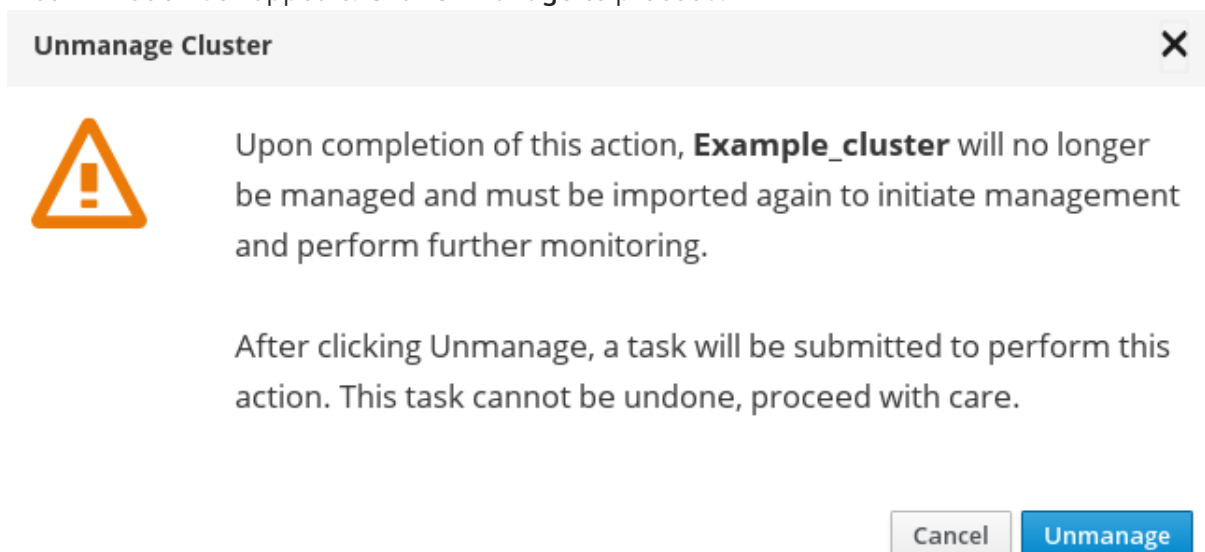
Unmanage a cluster from Web Administration

To unmanage a cluster:

1. Log in to the Web Administration interface and select the All Clusters view from drop-down.
2. Locate the cluster from the list of imported clusters. At the far right of the cluster row, click the vertical ellipsis and then click **Unmanage**.



3. A confirmation box appears. Click **Unmanage** to proceed.



4. The unmanage cluster task is submitted. To view the task progress, click **View Task Progress**.

Unmanage Cluster



An Unmanage cluster task is submitted to unmanage **Example_cluster**. You will be notified when processing is complete.

[View Task Progress](#)

- The status shows **Completed** after the unmanage task is successfully completed.

RED HAT GLUSTER STORAGE WEB ADMINISTRATION

Select a clust ▾

[All Clusters](#) » [Tasks](#) » UnmanageCluster: 0c6e2c22-a581-4175-a331-afccb03e507f

Time Submitted: 18 Jun 2018 17:14:17

Status: Completed

**NOTE**

After a cluster is unmanaged, it will not reappear immediately for reimport in the Web Administration environment. When the cluster is unmanaged, the `tendrl-node-agent` service undergoes synchronization with the nodes in the cluster. Based on the number of nodes in the cluster, expect a delay ranging from 60 seconds to a few minutes for the unmanaged cluster to reappear in the Web Administration environment.

4.1. GRAPHITE AND CARBON DATA ARCHIVAL

Graphite and Carbon Data Archival

When the cluster unmanage process is initiated, a concurrent process of cluster metrics data archival also begins. Once the cluster unmanage operation is successfully concluded, the Graphite and Carbon metrics data is archived to the following path on the Web Administration server:

```
/var/lib/carbon/whisper/tendrl/archive/clusters
```

The size of the archived data will depend on the size of the cluster that is unmanaged. You can either store the archived data or delete it to free some disk space.

4.2. GRAPHITE DATABASE METRICS RETENTION

The Graphite web service is a repository of telemetry data collected using `collectd`.

Changing metrics retention period

By default, Graphite retains the cluster metrics for a period of 180 days. To change the default metrics retention period for a given cluster, follow these steps:

1. Unmanage the cluster by following the [Unmanage Cluster](#) procedure outlined in the Red Hat Gluster Storage Web Administration Monitoring Guide.
2. After the cluster is unmanaged, stop the **carbon-cache** service on the Web Administration server:

```
# systemctl stop carbon-cache
```

3. Access the **storage-schemas.conf** file at:

```
/etc/tendr/monitoring-integration/storage-schemas.conf
```

4. Change the retention period in the **retentions** parameter under **[tendr]** section and save the changes.

Example:

If you want to change the retention value to 90 days, set the value in the **retentions** parameter:

```
[tendr]
pattern = ^tendr\.
retentions = 60s:90d
```

Where:

retentions = Each datapoint represents 60 seconds, and we want to keep enough datapoints so that they add up to 90 days of data.

5. Start the **carbon-cache** service on the Web Administration server:

```
# systemctl start carbon-cache
```

6. Import the cluster again. For instructions, see the [Import Cluster](#) chapter in the Red Hat Gluster Storage Web Administration Monitoring Guide.

CHAPTER 5. WEB ADMINISTRATION LOG LEVELS

In Web Administration, the log messages of different components and services are logged by **tendrl-node-agent**. It receives all the log messages from other components via a socket and logs the messages using **syslog**.

tendrl-node-agent logs messages for the following Web Administration components:

Server-side Web Administration components

- **tendrl-monitoring-integration**
- **tendrl-notifier**

Storage node-side Web Administration components

- **tendrl-gluster-integration**

These are the available log levels for Web Administration:

1. **DEBUG**
2. **INFO**
3. **WARNING**
4. **ERROR** (Default)
5. **CRITICAL**



NOTE

The log levels are case-sensitive. Input the log levels in uppercase format to avoid any error messages.

To change the default log level (**ERROR**), follow these steps:

1. Open the logging configuration yaml file in an editor:

```
| /etc/tendrl/node-agent/node-agent_logging.yaml
```

2. Change the log level by modifying all instances of **level** variables under **handlers** and **root**:

```
| handlers:
  console:
    class: logging.StreamHandler
    level: ERROR 1
    stream: ext://sys.stdout

  info_file_handler:
    class: logging.handlers.SysLogHandler
    facility: local5
    address: /dev/log
    level: ERROR 2
```

```
root:  
level: ERROR 3  
handlers: [console, info_file_handler]
```

1 Enter new log level

2 Enter new log level

3 Enter new log level

3. Restart the **tendrl-node-agent** service by:

```
# service tendrl-node-agent restart
```

CHAPTER 6. MONITORING AND METRICS

Gluster Web Administration provides deep metrics and visualization of Gluster clusters, the physical server nodes and the storage elements (disks) through the Grafana open-source monitoring platform.


CHAPTER 7. MONITORING DASHBOARD AND CONCEPTS

The Monitoring Dashboard provides high level visual information on health, performance and utilization of cluster wide resources.

7.1. DASHBOARD SELECTOR

The Dashboard Selector is the primary navigation tool to move between different dashboards.

Figure 7.1. Dashboard Selector




7.2. DASHBOARD PANELS

The Dashboard is composed of individual visualization blocks displaying different metrics and statistics termed as Panels. The panels exhibit different data based on the current status of the cluster component. Panels can be dragged and dropped and rearranged on the Dashboard.

The following are the types of panels available to visualize monitoring data:

1. **Single Status Panel:** The status panel displays the aggregated value of a series in e a single number data. For example, the Health, volume, snapshots are Singlestat panels.

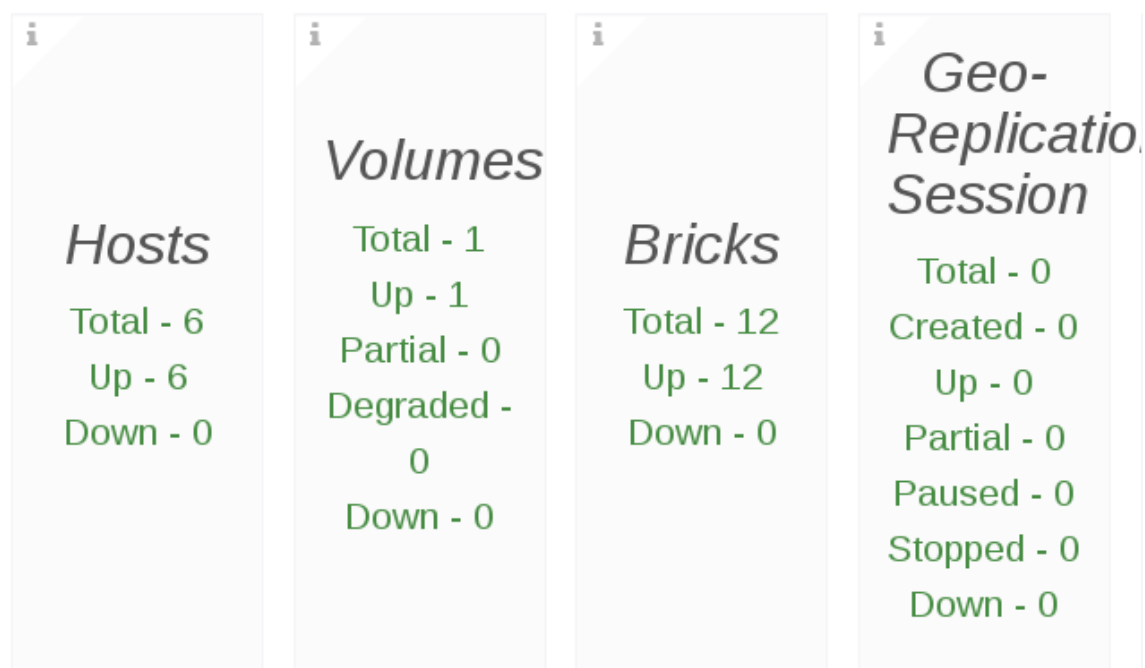
Figure 7.2. Single Status Panel Example




2. **Multiple Status Panel:** The Multiple status panel displays multiple values of a data source. The multiple status panel displays:

- The severity of the component
- If the component is disabled
- Extra data in the panel about the component

Figure 7.3. Multiple Status Panel Example




**NOTE**

If a multiple status panel displays **Invalid Number** as a status for a component, it indicates there is no data available to display from the time series database.

3. **Graph Panel:** The Graph panel allows to visualize unrestrained amounts of metrics. The Connection Trend and the Throughput Trend are examples of Graph panel.

Figure 7.4. Graph Panel Example



4. **Table Panel:** The Graph panel allows to visualize unrestrained amounts of metrics. The Connection Trend and the Throughput Trend are examples of Graph panel.

Figure 7.5. Table Panel Example

CPU Utilization by Host	
Host Name	Utilization ▾
tendr-usm1-gl3_example_org	7.8%
tendr-usm1-gl1_example_org	6.5%
tendr-usm1-gl2_example_org	6.5%
tendr-usm1-gl4_example_org	5.5%
tendr-usm1-gl6_example_org	4.7%

7.3. DASHBOARD ROWS

A row is a logical divider in a given Dashboard. The panels of the dashboard are arranged and organized in rows to give a streamlined look and visual.

7.4. DASHBOARD COLOR CODES

The Dashboard panels text displays the following color codes to represent health status information:

- **Green:** Healthy
- **Orange:** Degraded
- **Red:** Unhealthy, Down, or Unavailable

CHAPTER 8. MONITORING DASHBOARD FEATURES

8.1. DASHBOARD SEARCH

You can search the available dashboards by the dashboard name. The available filters to search a particular dashboard are starred and tags. The dashboard search functionality is accessed through the dashboard selector, located at the top of the Grafana interface.


Figure 8.1. Dashboard Search



8.2. DASHBOARD TIME RANGE

The Grafana interface provides time range management of the the data being visualized. You can change the time range for a graph to view the data at different points in time

At the top right, you can access the master Dashboard time picker. It shows the currently selected time range and the refresh interval.



Clicking the master Dashboard time picker, toggles a menu for time range controls.

Time range

From:

now-1h



To:

now



Refreshing every:

10s

Apply

Quick ranges

Last 2 days

Last 7 days

Last 30 days

Last 90 days

Last 6 months

Last 1 year

Last 2 years

Last 5 years

Yesterday

Day before yesterday

This day last week

Previous week

Previous month

Previous year

Today

Today so far

This week

This week so far

This month

This month so far

This year

This year so far

Last 5 minutes

Last 15 minutes

Last 30 minutes

Last 1 hour

Last 3 hours

Last 6 hours

Last 12 hours

Last 24 hours

Time range

The time range filter allows to mix both explicit and relative time ranges. The explicit time range format is **YYYY-MM-DD HH:MM:SS**.

Quick Range

Quick ranges are preset values to choose a relative time.

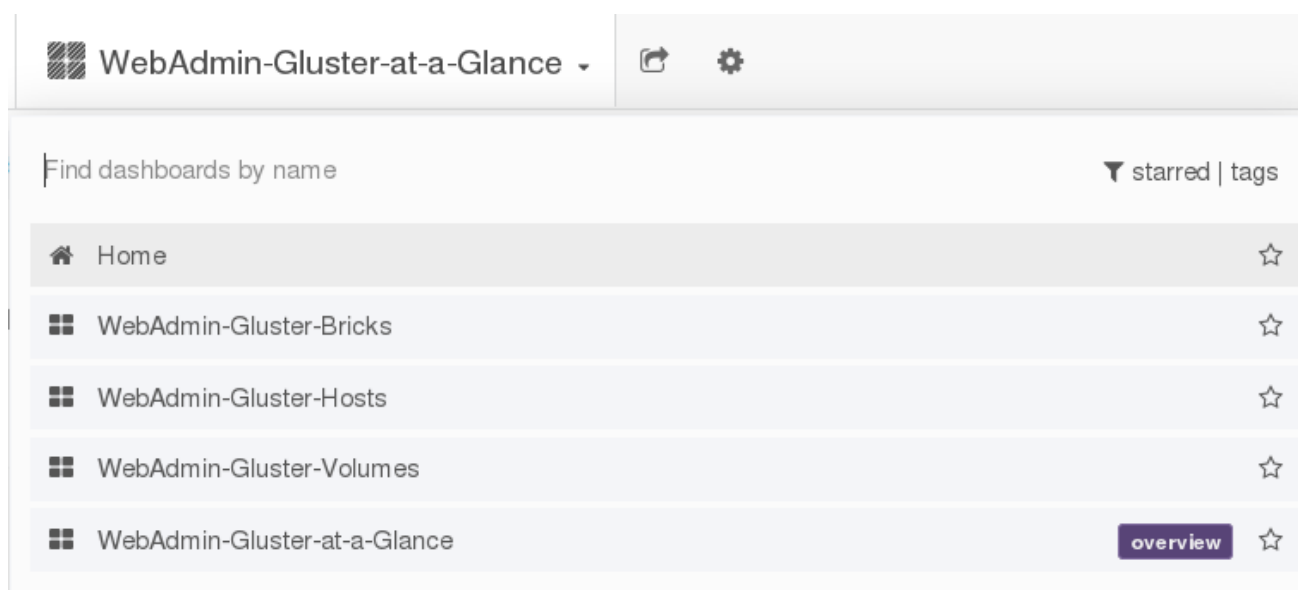
Refreshing every

When enabled, auto-refresh will reload the dashboard at the specified time range.

8.3. DASHBOARD SHARING

The Dashboard Selector is the primary navigation tool to move between different dashboards.

Figure 8.2. Dashboard Selector



CHAPTER 9. MONITORING DASHBOARD NAVIGATION

To access the Monitoring Dashboard, follow these steps:

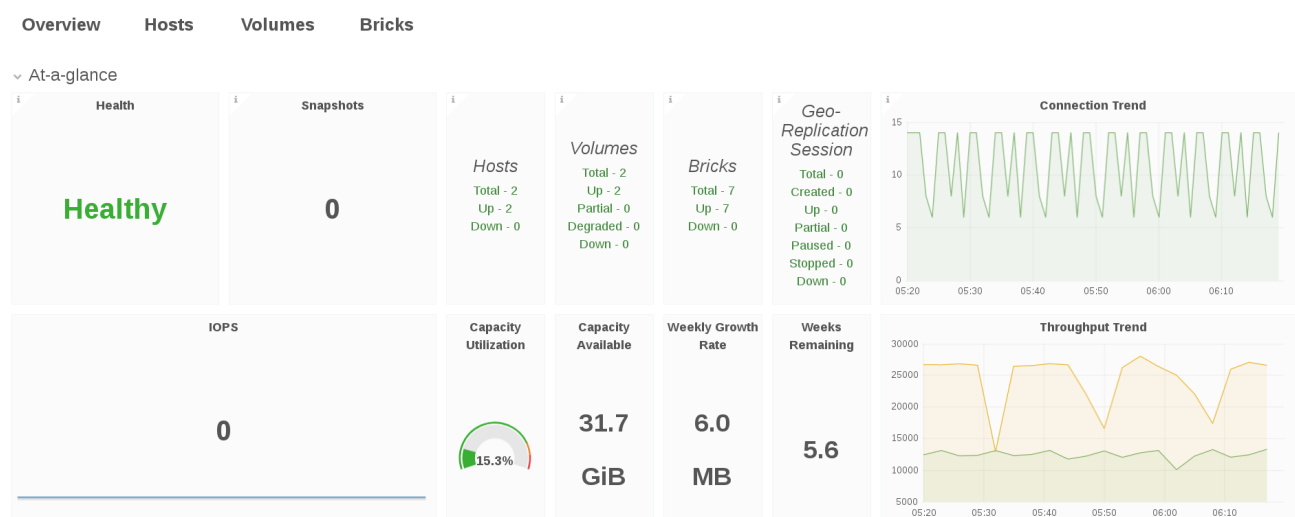
1. Log in to the Web Administration interface at <http://web-admin-server.example.com>.
2. In the default Cluster view, locate the cluster and click **Launch Dashboard**.
3. The Cluster dashboard showing the aggregated metrics view is opened in a new window.

9.1. CLUSTER VIEW DASHBOARD

The Cluster view dashboard allows the Gluster Administrator to:

- View at-a-glance information about the Gluster cluster that includes health and status information, key performance indicators such as IOPS, throughput, etc, and alerts that can highlight attention to potential issues in the cluster, host, volume, and brick.
- Compare a metric such as IOPS, CPU, Memory, Network Load across hosts within the cluster.
- Compare utilization across bricks within a volume, for example, IOPS, capacity, etc.

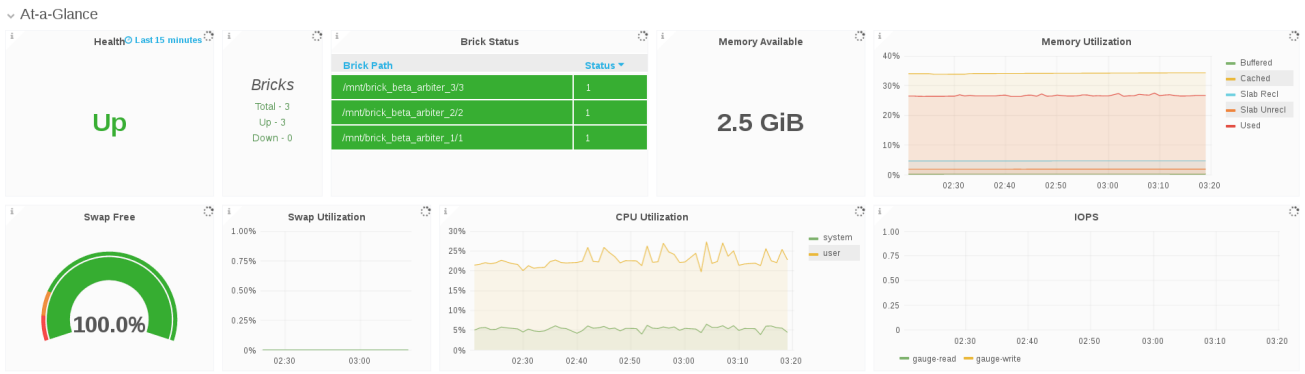
Figure 9.1. Cluster View Dashboard



9.2. HOSTS VIEW DASHBOARD

The Host view Dashboard allows the Gluster Administrator to:

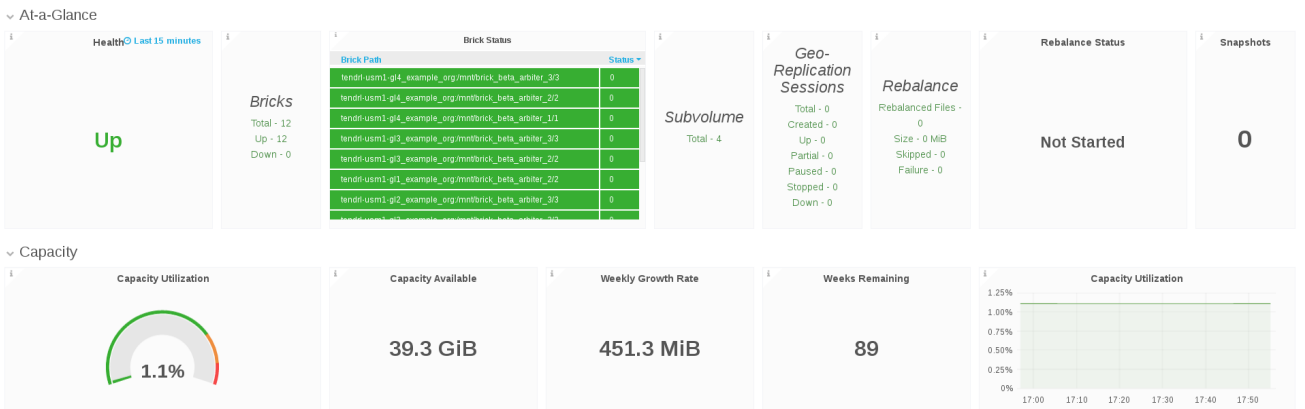
- View at-a-glance information about the Gluster host that includes health and status information, key performance indicators such as IOPS, throughput, etc and alerts that highlights attention to potential issues in the host, volume, brick, and disk.
- Compare one or more metrics such as IOPS, CPU, Memory, Network Load across bricks within the host.
- Compare utilization such as IOPS, capacity, etc across bricks within a host.



9.3. VOLUME VIEW DASHBOARD

The Volume view Dashboard allows the Gluster Administrator to:

- View at-a-glance information about the Gluster Volume that includes health and status information, key performance indicators that highlights attention to potential issues in the volume, brick, and disk.



9.4. BRICK VIEW DASHBOARD

The Brick view dashboard allows the Gluster Administrator to:

- View at-a-glance information about the Gluster brick that includes health and status information, key performance indicators such as IOPS, throughput, latency, etc and alerts that can highlight attention to potential issues in the brick and underlying disks.
- Look at performance by brick to address diagnosing of RAID 6 disk failure/rebuild/degradation poor performance on one brick.



CHAPTER 10. MONITORING CLUSTER METRICS

10.1. CLUSTER LEVEL DASHBOARD


This is the default dashboard of the Monitoring interface that shows the overview of the selected cluster.

10.1.1. Monitoring and Viewing Cluster Health

To monitor the Cluster health status and the metrics associated with it, view the panels in the Cluster Dashboard. For detailed panel descriptions and health indicators, see **Table 7.1. Cluster Health Panel Descriptions**.

10.1.1.1. Health and Snapshots

The **Health** panel displays the overall health of the selected cluster and the **Snapshots** panel shows the active number of snapshots.



10.1.1.2. Hosts, Volumes and Bricks

The Hosts, Volumes, and Bricks panels displays status information. The following is an example screen displaying the respective status information.



- **Hosts:** In total, there are 3 online Hosts
- **Volumes:** In total, there are 9 Volumes
- **Bricks:** In total, there are 44 Bricks

10.1.1.3. Geo-Replication Session

The Geo-Replication Session panel displays geo-replication session information from a given cluster, including the total number of geo-replication session and a count of geo-replication sessions by status.



10.1.1.4. Health Panel Descriptions

The following table lists the Panels and the descriptions.

Table 10.1. Cluster Health Panel Descriptions


Panel	Description	Health Indicator
Health	The Health panel displays the overall health of the selected cluster, which is either Healthy or Unhealthy	Green: Healthy Red: Unhealthy Orange: Degraded
Snapshots	The Snapshots panel displays the count of the active snapshots	
Hosts	The Hosts panel displays host status information including the total number of hosts and a count of hosts by status	
Volume	The Volumes panel displays volume status information for the selected cluster, including the total number of volumes and a count of volumes by status	
Bricks	The Bricks panel displays brick status information for the selected cluster, including the total number of bricks in the cluster, and a count of bricks by status	
Geo-Replication Session	The Geo-Replication Session panel displays geo-replication session information from a given cluster, including the total number of geo-replication session and a count of geo-replication sessions by status	

10.1.2. Monitoring and Viewing Cluster Performance

Cluster performance metrics can be monitored by the data displayed in the following panels.

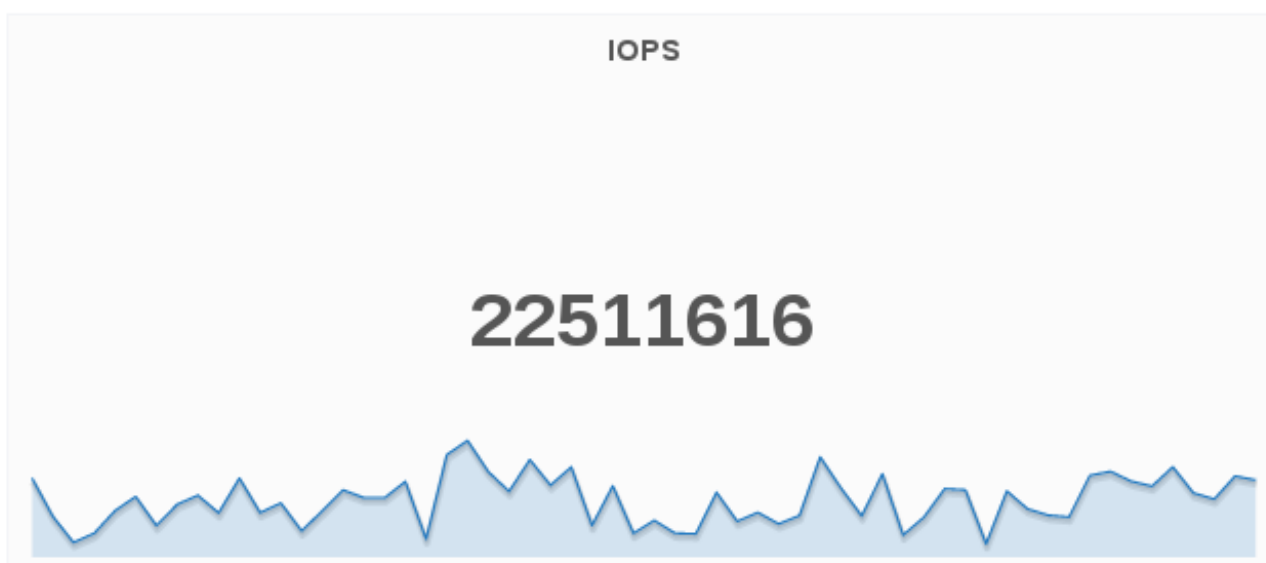
Connection Trend

The Connection Trend panel displays the total number of client connections to bricks in the volumes for the selected cluster over a period of time. Typical statistics may look like this:



IOPS

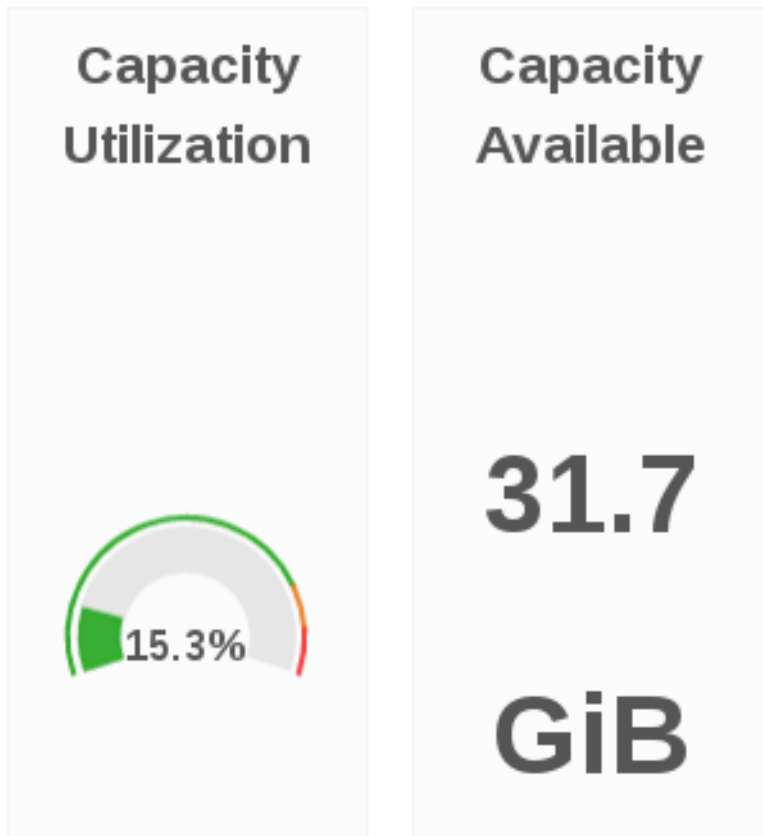
The IOPS panel displays IOPS for the selected cluster over a period of time. IOPS is based on the aggregated brick level read and write operations collected using *gluster volume profile* info.



Capacity Utilization and Capacity Available


The **Capacity Utilization** panel displays the capacity utilized across all volumes for the selected cluster.

The **Capacity Available** panel displays the available capacity across all volumes for the selected cluster.




Weekly Growth Rate

The Weekly Growth Rate panel displays the forecasted weekly growth rate for capacity utilization computed based on daily capacity utilization.




Weeks Remaining

The Weeks Remaining panel displays the estimated time remaining in weeks till volumes reach full capacity based on the forecasted Weekly Growth Rate.



Throughput Trend

The Throughput Trend panel displays the network throughput for the selected cluster over a period of time.



10.1.3. Top Consumers

The Top Consumers panels displays the highest capacity utilization by the cluster resources.

To view the top consumers of the cluster:

1. In the Cluster level dashboard, at the bottom, click **Top Consumers** to expand the menu.

> **Top Consumers**

> **Status**

Top 5 Utilization By Bricks

The Top 5 Utilization By Bricks panel displays the bricks with the highest capacity utilization.

Top 5 Utilization by Bricks	
Bricks	Utilization ▼
tendrl-usm1-gl2_example_org:/mnt/brick_beta_arbiter_2/2	0.1%
tendrl-usm1-gl2_example_org:/mnt/brick_beta_arbiter_1/1	0.1%
tendrl-usm1-gl1_example_org:/mnt/brick_beta_arbiter_3/3	0.1%
tendrl-usm1-gl1_example_org:/mnt/brick_beta_arbiter_2/2	0.1%
tendrl-usm1-gl1_example_org:/mnt/brick_beta_arbiter_1/1	0.1%

Top 5 Utilization by Volume

The Top 5 Utilization By Volumes panel displays the volumes with the highest capacity utilization.

Top 5 Utilization by Volume	
Volume Name	Utilization ▼
volume_beta_arbiter_2_plus_1x2	1.1%
volume_alpha_distrep_6x2	1.1%

CPU Utilization by Host

The CPU Utilization by Host panel displays the CPU utilization of each node in the cluster.

CPU Utilization by Host	
Host Name	Utilization ▾
tendrl-usm1-gl4_example_org	12.0%
tendrl-usm1-gl3_example_org	7.8%
tendrl-usm1-gl1_example_org	5.0%
tendrl-usm1-gl2_example_org	4.7%


Memory Utilization By Host

The Memory Utilization by Hosts panel displays memory utilization of each node in the cluster.

Memory Utilization by Host	
Host Name	Utilization ▾
tendrl-usm1-gl1_example_org	6.2%
tendrl-usm1-gl2_example_org	6.1%
tendrl-usm1-gl3_example_org	6.0%
tendrl-usm1-gl4_example_org	6.0%
tendrl-usm1-gl5_example_org	5.6%

Ping Latency Trend


The Ping Latency Trend panel displays the ping latency for each host in a given cluster.



10.1.4. Monitoring and Viewing Cluster Status

To view the status of the overall cluster:

1. In the Cluster level dashboard, at the bottom, click **Status** to expand the menu.



2. The Volume, Host, and Brick status are displayed in the panels.

Volume Status		Host Status		Brick Status	
Volume Name	Status	Host Name	Status	Brick Name	Status
volume_beta_arbiter_2_plus_1x2	0	tendri-usr1-gl6_example_org	0	tendri-usr1-gl4_example_org/mnt/brick_beta_arbiter_3/3	0
		tendri-usr1-gl5_example_org	0	tendri-usr1-gl4_example_org/mnt/brick_beta_arbiter_2/2	0
		tendri-usr1-gl4_example_org	0	tendri-usr1-gl4_example_org/mnt/brick_beta_arbiter_1/1	0
		tendri-usr1-gl3_example_org	0	tendri-usr1-gl3_example_org/mnt/brick_beta_arbiter_3/3	0
		tendri-usr1-gl2_example_org	0	tendri-usr1-gl3_example_org/mnt/brick_beta_arbiter_2/2	0
		tendri-usr1-gl1_example_org	0	tendri-usr1-gl1_example_org/mnt/brick_beta_arbiter_1/1	0

Volume Status

The Volume Status panel displays the status code of each volume for the selected cluster.

Volume Status	
Volume Name	Status
volume_gama_disperse_4_plus_2x2	0

The volume status is displayed in numerals and colors. The following are the corresponding status of the numerals.

- 0 = Up
- 3 = Up (Degraded)
- 4 = Up (Partial)
- 5 = Unknown
- 8 = Down

Host Status

The Host Status panel displays the status code of each host for the selected cluster.

Host Status	
Host Name	Status ▾
tendrl-usm1-gl6_example_org	0
tendrl-usm1-gl5_example_org	0
tendrl-usm1-gl4_example_org	0
tendrl-usm1-gl3_example_org	0
tendrl-usm1-gl2_example_org	0
tendrl-usm1-gl1_example_org	0

The Host status is displayed in numeric codes:

- 0 = Up
- 8 = Down

Brick Status

The Brick Status panel displays the status code of each brick for the selected cluster.

Brick Status	
Brick Name	Status ▾
tendrl-usm1-gl4_example_org:/mnt/brick_beta_arbiter_3/3	0
tendrl-usm1-gl4_example_org:/mnt/brick_beta_arbiter_2/2	0
tendrl-usm1-gl4_example_org:/mnt/brick_beta_arbiter_1/1	0
tendrl-usm1-gl3_example_org:/mnt/brick_beta_arbiter_3/3	0
tendrl-usm1-gl3_example_org:/mnt/brick_beta_arbiter_2/2	0
tendrl-usm1-gl1_example_org:/mnt/brick_beta_arbiter_2/2	0

The Brick status is displayed in numeric codes:

- 1 = Started
- 10 = Stopped


10.2. HOST LEVEL DASHBOARD

10.2.1. Monitoring and Viewing Health and Status

To monitor the Cluster Hosts status and the metrics associated with it, navigate to the Hosts Level Dashboard and view the panels.

Health

The Health panel displays the overall health for a given host.



Bricks and Bricks Status

The Bricks panel displays brick status information for a given host, including the total number of bricks in the host, and a count of bricks by status.



The Brick Status panel displays the status code of each brick for a given host.

Brick Status	
Brick Path	Status ▾
/root/bricks/v3	1
/root/bricks/v22	1
/root/bricks/v2	1


- 1 = Started
- 10 = Stopped

10.2.2. Monitoring and Viewing Performance

10.2.2.1. Memory and CPU Utilization

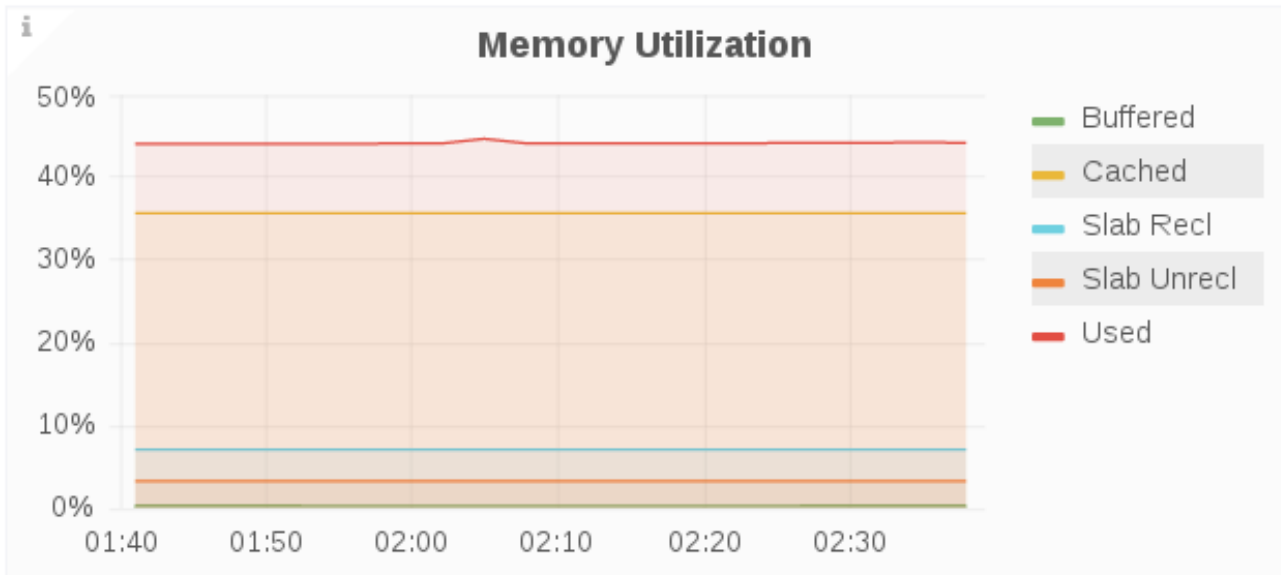
Memory Available

The Memory Available panel displays the sum of memory free and memory cached.



Memory Utilization

The Memory Utilization panel displays memory utilization percentage for a given host that includes buffers and caches used by the kernel over a period of time.



- **Buffered:** Amount of memory used for buffering, mostly for I/O operations
- **Cached:** Memory used for caching disk data for reads, memory-mapped files or tmpfs data
- **Slab Rec:** Amount of reclaimable memory used for slab kernel allocations
- **Slab Unrecl:** Amount of unreclaimable memory used for slab kernel allocations
- **Used:** Amount of memory used, calculated as Total - Free (Unused Memory) - Buffered - Cache
- **Total:** Total memory used


Swap Free

The Swap Free panel displays the available swap space in percent for a given host.




Swap Utilization

The Swap Utilization panel displays the used swap space in percent for a given host.




CPU Utilization

The CPU utilization panel displays the CPU utilization for a given host over a period of time.



IOPS


The IOPS panel displays IOPS for a given host over a period of time. IOPS is based on the aggregated brick level read and write operations.



10.2.2.2. Capacity and Disk Load


Total Brick Capacity Utilization Trend

The Total Brick Capacity Utilization Trend panel displays the capacity utilization for all bricks on a given for a period of time.




Total Brick Capacity Utilization

The Total Brick Capacity Utilization panel displays the current percent capacity utilization for a given host.



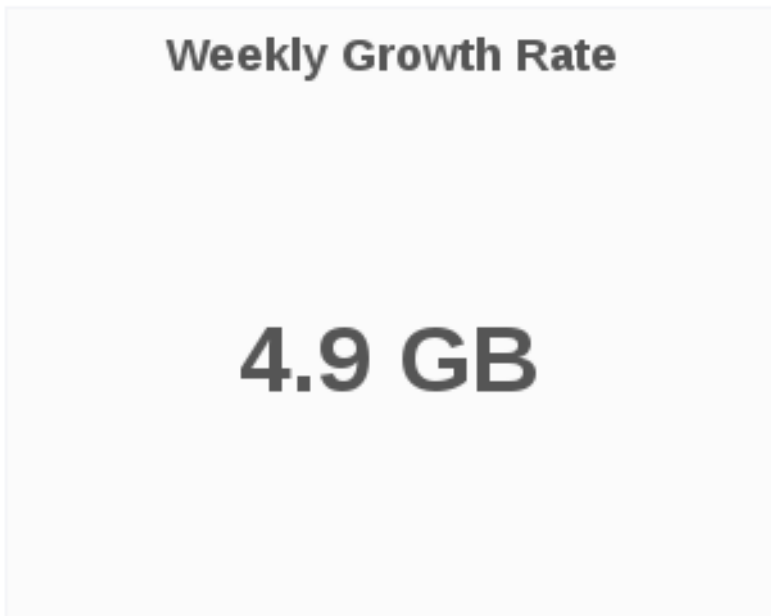
Total Brick Capacity Available

The Total Brick Capacity Available panel displays the current available capacity for a given host.




Weekly Growth Rate

The Weekly Growth Rate panel displays the forecasted weekly growth rate for capacity utilization computed based on daily capacity utilization.



Weeks Remaining

The Weeks Remaining panel displays the estimated time remaining in weeks till host capacity reaches full capacity based on the forecasted Weekly Growth Rate.



Brick Utilization

The Brick Utilization panel displays the utilization of each brick for a given host.

Brick Utilization	
Brick Path	Utilization ▼
/mnt/brick_gama_disperse_2/2	0.2%
/mnt/brick_gama_disperse_1/1	0.2%

Brick Capacity

The Brick Capacity panel displays the total capacity of each brick for a given host.

Brick Capacity	
Brick Path	Capacity ▾
/mnt/brick_gama_disperse_2/2	19.9 GiB
/mnt/brick_gama_disperse_1/1	19.9 GiB


Brick Capacity Used

The Brick Capacity Used panel displays the used capacity of each brick for a given host.

Brick Capacity Used	
Brick Path	Capacity ▾
/mnt/brick_gama_disperse_2/2	33.0 MiB
/mnt/brick_gama_disperse_1/1	33.0 MiB


Disk Load

The Disk Load panel shows the host's aggregated read and writes from/to disks over a period of time.




Disk Operation

The Disk Operations panel shows the host's aggregated read and writes disk operations over a period of time.



Disk IO


The Disk IO panel shows the host's aggregated I/O time over a period of time.



10.2.2.3. Network


Throughput

The Throughput panel displays the network throughput for a given host over a period of time.




Dropped Packets Per Second

The Dropped Packets Per Second panel displays dropped network packets for the host over a period of time. Typically, dropped packets indicates network congestion, for example, the queue on the switch port your host is connected to is full and packets are dropped because it cannot transmit data fast enough.



Errors Per Second

The Errors Per Second panel displays network errors for a given host over a period of time. Typically, the errors indicate issues that occurred while transmitting packets due to carrier errors (duplex mismatch, faulty cable), fifo errors, heartbeat errors, and window errors, CRC errors too short frames, and/or too long frames. In short, errors typically result from faulty hardware, and/or speed mismatch.



10.2.3. Host Dashboard Metric Units

The following table shows the metrics and their corresponding measurement units.

Table 10.2. Host Dashboard Metric Units

Metrics	Units
Memory Available	Megabyte/Gigabyte/Terabyte
Memory Utilization	Percentage %
Swap free	Percentage %
Swap Utilization	Percentage %
CPU Utilization	Percentage %
Total Brick Capacity Utilization	Percentage %
Total Brick Capacity	MB/GB/TB
Weekly Growth Rate	MB/GB/TB
Disk Load	kbps
Disk IO	millisecond ms
Network Throughput	kbps

10.3. VOLUME LEVEL DASHBOARD

The Volume view dashboard allows the Gluster Administrator to:

- View at-a-glance information about the Gluster volume that includes health and status information, key performance indicators such as IOPS, throughput, etc, and alerts that can highlight attention to potential issues in the volume, brick, and disk.
- Compare 1 or more metrics such as IOPS, CPU, Memory, Network Load across bricks within the volume.
- Compare utilization such as IOPS, capacity, etc, across bricks within a volume.
- View performance metrics by brick (within a volume) to address diagnosing of failure, rebuild, degradation, and poor performance on one brick.

When all the Gluster storage nodes are shut down or offline, Time to live (TTL) will delete the volume details from etcd as per the TTL value measured in seconds. The TTL value for volumes is set based on the number of volumes and bricks in the system. The formula to calculate the TTL value to delete volume details is:

Time to Live (seconds) = synchronization interval (60 seconds) + number of volumes * 20 + number of bricks * 10 + 160.

In Web Administration environment

- Cluster will show status as unhealthy and all hosts will be marked as down
- No display of Volumes and Bricks
- The Events view will reflect the relevant status

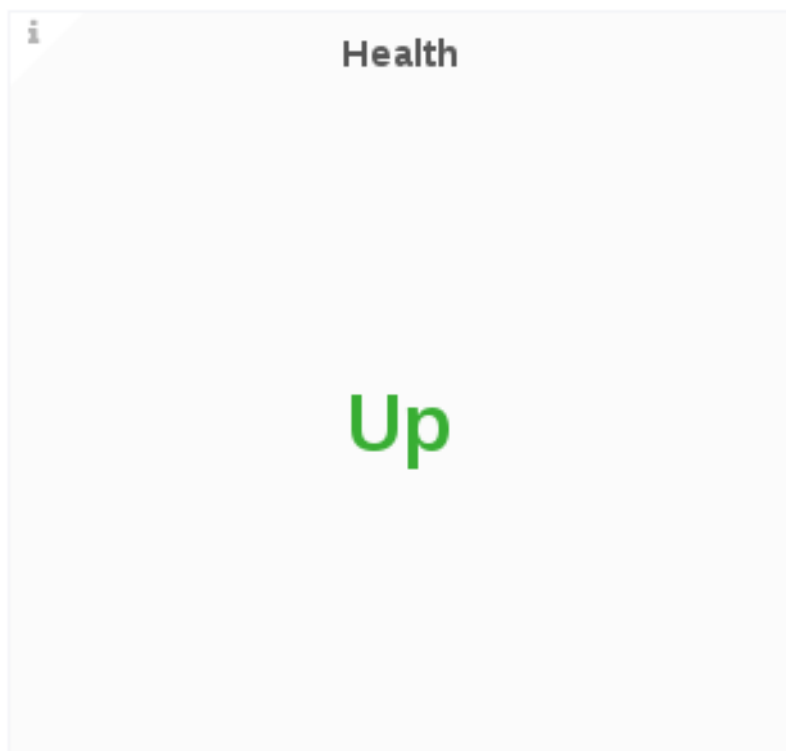
In Grafana Dashboard

- In Cluster level Dashboard, the Host, Volumes, and Bricks panels reflects the relevant updated counts with status.
- In Cluster, Volume, and Brick level dashboards, some panels will be marked as N/A, indicating no data is available.

10.3.1. Monitoring and Viewing Health


Health

The Health panel displays the overall health for a given volume.



Snapshots

The Snapshots panel displays the count of active snapshots for the selected cluster.



Brick Status

The Brick Status panel displays the status code of each brick for a given volume.

Brick Status	
Brick Path	Status
tendrl-usm1- gl4_example_org:/mnt/brick_beta_arbiter_3/3	0
tendrl-usm1- gl4_example_org:/mnt/brick_beta_arbiter_2/2	0
tendrl-usm1- gl4_example_org:/mnt/brick_beta_arbiter_1/1	0
tendrl-usm1- gl3_example_org:/mnt/brick_beta_arbiter_3/3	0
tendrl-usm1-	0

- 1 = Started
- 10 = Stopped

Bricks

The Bricks panel displays brick status information for a given volume, including the total number of bricks in the volume, and a count of bricks by status.



Subvolumes

The Subvolumes panel displays subvolume status information for a given volume.

Subvolume

Total - 2

Geo-Replication Sessions

The Geo-Replication Session panel displays geo-replication session information from a given volumes, including the total number of geo-replication session and a count of geo-replication sessions by status.

Geo-Replication Sessions

Total - 0

Created - 0

Up - 0

Partial - 0

Paused - 0

Stopped - 0

Down - 0

Rebalance

The Rebalance panel displays rebalance progress information for a given volume, which is applicable when rebalancing is underway.

Rebalance

Rebalanced

Files - 0

MB Size - 0

Skipped - 0

Failure - 0

Rebalance Status:

The Rebalance Status panel displays the status of rebalancing for a given volume, which is applicable when rebalancing is underway.

**Rebalance
Status**


Not

Started

10.3.2. Monitoring and Viewing Performance


Capacity Utilization

The Capacity Utilization panel displays the used capacity for a given volume.




Capacity Available

The Capacity Available panel displays the available capacity for a given volume.




Weekly Growth Rate

The Weekly Growth Rate panel displays the forecasted weekly growth rate for capacity utilization computed based on daily capacity utilization.




Weeks Remaining

The Weeks Remaining panel displays the estimated time remaining in weeks till volume reaches full capacity based on the forecasted Weekly Growth Rate.




Capacity Utilization Trend

The Capacity Utilization Trend panel displays the volume capacity utilization over a period of time.



Inode Utilization

The Inode Utilization panel displays inodes used for bricks in the volume over a period of time.



Inode Available

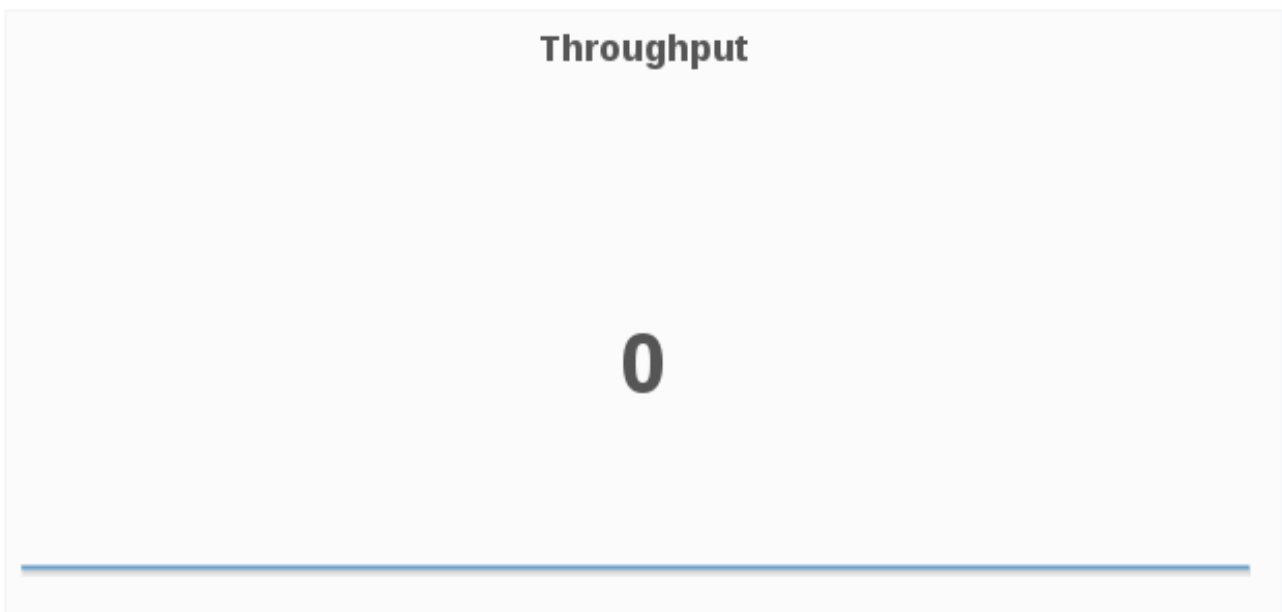
The Inode Available panel displays inodes free for bricks in the volume.

Inode Available

20287947


Throughput

The Throughput panel displays volume throughput based on brick-level read and write operations fetched using *gluster volume profile*.




LVM Thin Pool Metadata %

The LVM Thin Pool Metadata % panel displays the utilization of LVM thin pool metadata for a given volume. Monitoring the utilization of LVM thin pool metadata and data usage is important to ensure they do not run out of space. If the data space is exhausted, I/O operations are either queued or failing based on the configuration. If metadata space is exhausted, you will observe error I/O's until the LVM pool is taken offline and repair is performed to fix potential inconsistencies. Moreover, due to the metadata transaction being aborted and the pool doing caching there might be uncommitted (to disk) I/O operations that were acknowledged to the upper storage layers (file system) so those layers will need to have checks/repairs performed as well.



LVM Thin Pool Data Usage %

The LVM Thin Pool Data Usage % panel displays the LVM thin pool data usage for a given volume. Monitoring the utilization of LVM thin pool metadata and data usage is important to ensure they do not run out of space. If the data space is exhausted, I/O operations are either queued or failing based on the configuration. If metadata space is exhausted, you will observe error I/O's until the LVM pool is taken offline and repair is performed to fix potential inconsistencies. Moreover, due to the metadata transaction being aborted and the pool doing caching there might be uncommitted (to disk) I/O operations that were acknowledged to the upper storage layers (file system) so those layers will need to have checks/repairs performed as well.



10.3.3. Monitoring File Operations

Top File Operations

The Top File Operations panel displays the top 5 FOP (file operations) with the highest % latency, wherein the % latency is the fraction of the FOP response time that is consumed by the FOP.

Top File Operations	
File Operations	% Latency ▾
LOOKUP	16.00
GETXATTR	9.00
READDIR	8.00
STATFS	4.00
OPENDIR	4.00

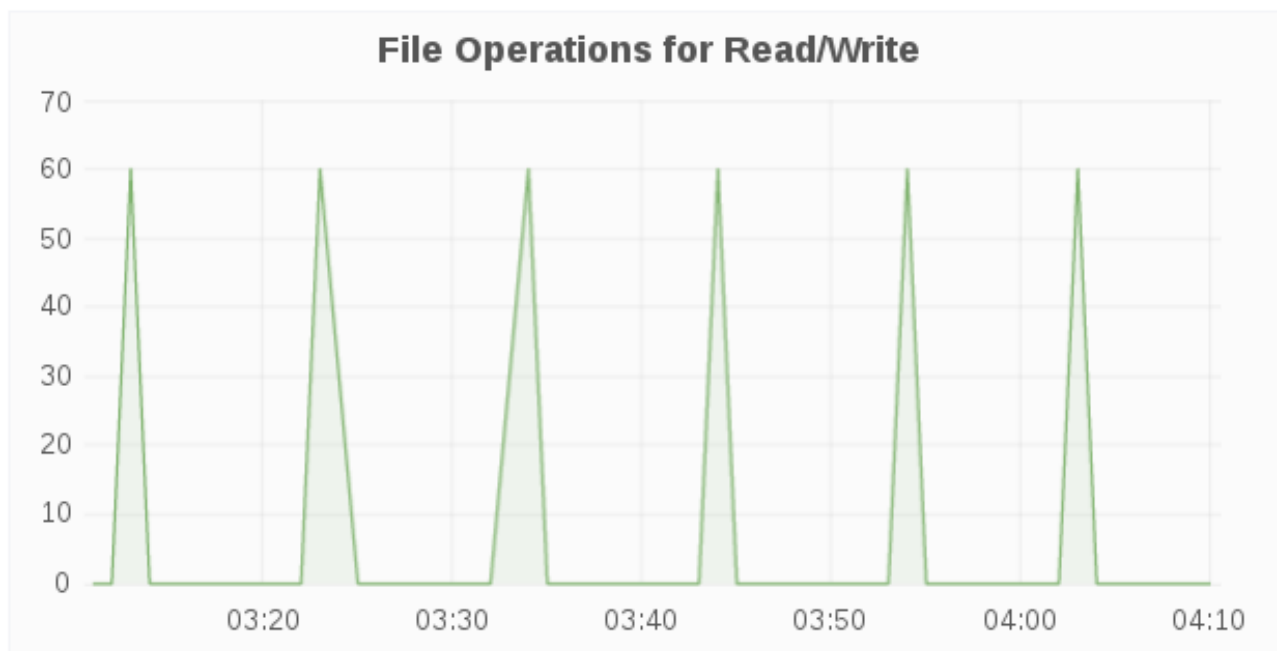
File Operations for Locks Trend

The File Operations for Locks Trend panel displays the average latency, maximum latency, call rate for each FOP for Locks over a period of time.




File Operations for Read/Write

The File Operations for Read/Write panel displays the average latency, maximum latency, call rate for each FOP for Read/Write Operations over a period of time.




File Operations for Inode Operations

The File Operations for Inode Operations panel displays the average latency, maximum latency, call rate for each FOP for Inode Operations over a period of time.



File Operations for Entry Operations

The File Operations for Entry Operations panel displays the average latency, maximum latency, call rate for each FOP for Entry Operations over a period of time.



10.3.4. Volume Dashboard Metric Units

The following table shows the metrics and their corresponding measurement units.


Table 10.3. Volume Dashboard Metric Units

Metrics	Units
Capacity Utilization	Percentage %
Capacity Available	Megabyte/Gigabyte/Terabyte
Weekly Growth Rate	Megabyte/Gigabyte/Terabyte
Capacity Utilization Trend	Percentage %
Inode Utilization	Percentage %
Lvm Thin Pool Metadata	Percentage %
Lvm Thin Pool Data Usage	Percentage %
File Operations for Locks Trend	MB/GB/TB
File Operations for Read/Write	K
File Operations for Inode Operation Trend	K
File Operations for Entry Operations	K

10.4. BRICK LEVEL DASHBOARD

10.4.1. Monitoring and Viewing Brick Status


The Status panel displays the status for a given brick.



10.4.2. Monitoring and Viewing Brick Performance


Capacity Utilization

The Capacity Utilization panel displays the percentage of capacity utilization for a given brick.




Capacity Available

The Capacity Available panel displays the available capacity for a given volume.




Capacity Utilization Trend

The Capacity Utilization Trend panel displays the brick capacity utilization over a period of time.



Weekly Growth Rate

The Weekly Growth Rate panel displays the forecasted weekly growth rate for capacity utilization computed based on daily capacity utilization.





Weeks Remaining

The Weeks Remaining panel displays the estimated time remaining in weeks till brick reaches full capacity based on the forecasted Weekly Growth Rate.



Healing

The Healing panel displays healing information for a given volume based on healinfo.




NOTE

The Healing panel will not show any data for volumes without replica.


IOPS

The IOPS panel displays IOPS for a brick over a period of time. IOPS is based on brick level read and write operations.




LVM Thin Pool Metadata %

The LVM Thin Pool Metadata % panel displays the utilization of LVM thin pool metadata for a given brick. Monitoring the utilization of LVM thin pool metadata and data usage is important to ensure they don't run out of space. If the data space is exhausted, I/O operations are either queued or failing based on the configuration. If metadata space is exhausted, you will observe error I/O's until the LVM pool is taken offline and repair is performed to fix potential inconsistencies. Moreover, due to the metadata transaction being aborted and the pool doing caching there might be uncommitted (to disk) I/O operations that were acknowledged to the upper storage layers (file system) so those layers will need to have checks/repairs performed as well.




LVM Thin Pool Data Usage %

The LVM Thin Pool Data Usage % panel displays the LVM thin pool data usage for a given brick. Monitoring the utilization of LVM thin pool metadata and data usage is important to ensure they don't run out of space. If the data space is exhausted, I/O operations are either queued or failing based on the configuration. If metadata space is exhausted, you will observe error I/O's until the LVM pool is taken offline and repair is performed to fix potential inconsistencies. Moreover, due to the metadata transaction being aborted and the pool doing caching there might be uncommitted (to disk) I/O operations that were acknowledged to the upper storage layers (file system) so those layers will need to have repairs performed as well.




Throughput

The Throughput panel displays brick-level read and write operations fetched using "gluster volume profile."



Latency

The Latency panel displays latency for a brick over a period of time. Latency is based on the average amount of time a brick spends doing a read or write operation.



10.4.3. Brick Dashboard Metric Units

The following table shows the metrics and their corresponding measurement units.

Table 10.4. Brick Dashboard Metric Units

Metrics	Units
Capacity Utilization	Percentage %
Capacity Available	Megabyte/Gigabyte/Terabyte
Weekly Growth Rate	Megabyte/Gigabyte/Terabyte
Capacity Utilization Trend	Percentage %
Inode Utilization	Percentage %
Lvm Thin Pool Metadata	Percentage %
Lvm Thin Pool Data Usage	Percentage %
Disk Throughput	Percentage %

CHAPTER 11. USERS AND ROLES ADMINISTRATION

11.1. USER ROLES

There are three user roles available for Web Administration.

1. **Admin:** The Admin role gives complete rights to the user to manage all Web Administration operations.
2. **Normal User:** The Normal User role authorizes the user to perform operations such as importing cluster and enabling or disabling volume profiling but restricts managing users and other administrative operations.
3. **Read-only User:** Read-only: The Read-only User role authorizes the user to only view and monitor cluster-wide metrics and readable data. The user can launch Grafana dashboards from the Web Administration interface but is restricted to perform any storage operations. This role is suited for users performing monitoring tasks.

11.2. CONFIGURING ROLES

To add and configure a new user, follow these steps:

1. Log In the Web Administration interface and in navigation pane, click **Admin > Users**.
2. The users list is displayed. To add a new user, click **Add** at the right-hand side.

User ID ▾ *Filter by Use.*
Add

User ID	Name	Role	Notification	Email	Actions
admin	Admin	Admin	Disabled	admin@tendrl.org	Edit ⋮

3. Enter the user information in the given fields. To enable or disable email notifications, toggle the **ON-OFF** button.





Add a new local user. All user information will be stored internally in Tendrl.

User ID	Normal
Name	John Smith
Email	jsmith@org.com
	Email Notifications ON
Password 👁
Confirm Password 👁

4. Select a **Role** from the available three roles and click **Save**.

- Select a Role
- Admin**
The Admin role permits the user to manage Tendrl and all storage operations.
 - Normal User**
The Normal user role permits the user to provision storage but cannot manage the Tendrl applications.
 - Read-only User**
The Read-only user role permits the user to access Tendrl in read-only mode, which means the user cannot actively manage storage. This role is ideal for users interested in monitoring the storage.


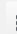

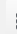
5. The new user is successfully created.

User ID	Name	Role	Notification	Email	Actions
admin	Admin	Admin	 Disabled	admin@tendrl.org	<input type="button" value="Edit"/> 
administrator	John Smith	Normal	 Enabled	jsmith@org.com	<input type="button" value="Edit"/> 

11.2.1. Editing Users

To edit an existing user:

1. Navigate to the user view by clicking **Admin > Users** from the interface navigation.
2. Locate the user to be edited and click **Edit** at the right-hand side.

User ID	Name	Role	Notification	Email	Actions
admin	Admin	Admin	 Disabled	admin@tendrl.org	<input type="button" value="Edit"/> 
administrator	John Smith	Normal	 Enabled	jsmith@org.com	<input type="button" value="Edit"/> 

3. Edit the required information and click **Save**.

User ID Normal

Name

Email Email Notifications ON

New Password

Confirm Password

11.2.2. Disabling Notifications and Deleting User

Enabling and Disabling Notifications

To enable notifications:

1. Navigate to the user view by clicking **Admin > Users** from the interface navigation.

User ID	Name	Role	Notification	Email	Actions
admin	Admin	Admin	Disabled	admin@tendrl.org	<input type="button" value="Edit"/>
administrator	John Smith	Normal	Enabled	jsmith@org.com	<input type="button" value="Edit"/>

2. Click the vertical elipsis next to the Edit button and click **Disable Email Notification** from the callout menu.

Email	Actions
admin@tendrl.org	<input type="button" value="Edit"/>
jsmith@org.com	<input type="button" value="Edit"/>

Disable Email Notification
Delete User

3. Email notification is successfully disabled for the user.

Email notification is now disabled for Normal. ✕

Users

User ID ▾ *Filter by User ID*[Add](#)

User ID	Name	Role	Notification	Email	Actions
admin	Admin	Admin	Disabled	admin@tendr1.org	Edit ⋮
Normal	John Smith	Normal	Disabled	jsmith@org.com	Edit ⋮

Deleting User

To delete an existing user:

1. Navigate to the user view by clicking **Admin > Users** from the interface navigation.
2. Locate the user to be deleted and click the vertical elipsis next to the Edit button. A callout menu opens, click **Delete User**.

Email	Actions
admin@tendr1.org	Edit ⋮
jsmith@org.com	Edit ⋮

Disable Email Notification

Delete User

3. A confirmation box appears. Click **Delete**.

Delete User ✕

This action will permanently remove user **Normal**. Click Delete to continue.

[Cancel](#)[Delete](#)

CHAPTER 12. ALERTS AND NOTIFICATIONS

Alerts are current problems and critical conditions that occur in the system and notified to the user. The Grafana monitoring platform generates alerts based on severity levels.

You can configure alerts via SMTP and SNMP protocols. SMTP configuration will send email alerts to users that have email notifications enabled. SNMPv3 configuration will send SNMP trap alerts to the Alerts notifications drawer of the Web Administration environment.

12.1. TYPES OF ALERTS

The alerts triggered by the dashboard are classified in the following categories:

- **Status alerts** : Alerts arising when a cluster resource undergoes a change of state. For example, Healthy to Unhealthy.
- **Utilization alerts**: Alerts arising after a cluster resource exceed the set threshold and after it reverts to the normal state. For example, when the Host CPU utilization is breached, an alert is triggered notifying the user about the event.

12.2. LIST OF ALERTS

The list of Web Administration alerts are given in the tables below.

Status Alerts

Table 12.1. Status Alerts

Alert	System Resource(s)
volume status	Volume and Cluster
volume state	Volume and Cluster
brick status	Volume, Host, and Cluster
peer status	Cluster
rebalance status	Volume and Cluster
Geo-Replication status	Cluster
quorum of volume lost	Volume and Cluster
quorum of volume regained	Volume and Cluster
svc connected	Cluster
svc disconnected	Cluster

Alert	System Resource(s)
minimum number of bricks not up in EC subvolume	Volume and Cluster
minimum number of bricks up in EC subvolume	Volume and Cluster
afr quorum met for subvolume	Volume and Cluster
afr quorum fail for subvolume	Volume and Cluster
afr subvolume up	Volume and Cluster
afr subvolume down	Volume and Cluster

Utilization Alerts

Table 12.2. Utilization Alerts


Alert	System Resource
cpu utilization	Host
memory utilization	host
swap utilization	host
volume utilization	Volume and Cluster
brick utilization	Volume and Cluster

12.3. ALERTS NOTIFICATIONS DRAWER

Alerts drawer is a notification indicator embedded in the Web Administration interface to display the system wide alerts.

Accessing Alerts Drawer

1. To access the Alerts drawer, log in the Web Administration interface. In the default landing interface, locate and click on the interactive bell icon on the header bar at the top right-hand side.





2. The drawer is opened displaying the number of alerts generated.


⚙️ 🔔 👤

Alerts ✕

3 Alerts 📘 1 ⚠️ 2

 Cluster:ClusterA is unhealthy
28 Aug 2018 13:59:23

 Volume:volume_beta_arbiter_2_plus_1x2 is down
28 Aug 2018 13:59:22

 Peer tendrl-usm1-gl1.example.org in cluster ClusterA is Connected
28 Aug 2018 13:59:18

To filter alerts, click on the status icon at the right.



12.4. SMTP NOTIFICATIONS CONFIGURATION

Tendrl-ansible installs and configures tendrl-notifier. After the tendrl-notifier file is configured, configure SMTP email notifications:

1. Open the `/etc/tendrl/notifier/email.conf.yaml` file.
2. Update the parameters:

```
email_id = <The sender email id>
email_smtp_server = <The smtp server>
email_smtp_port = <The smtp port>
```

3. If the SMTP server supports only authenticated email, follow the template in the `/etc/tendrl/notifier/email_auth.conf.yaml` file and accordingly enable the following:


```
auth = <ssl/tls>
```

```
email_pass = <password corresponding to email_id for authenticating to smtp server>
```

- Restart the tendrl-notifier service:

```
systemctl restart tendrl-notifier
```

12.5. SNMPV3 NOTIFICATION CONFIGURATION

Configure SNMP

Tendrl-ansible installs and configures tendrl-notifier. After the tendrl-notifier file is configured, configure SNMPv3 trap notifications:

- Open the **tendrl-notifier** configuration file:

```
# cat /etc/tendrl/notifier/snmp.conf.yaml
```

- Update the parameters in the file for v3 trap alerts:

```
For v3_endpoint:
# For more hosts you can add more entry with endpoint2, endpoint3, etc
endpoint1:

# Name or IP address of the remote SNMP host.
  host_ip: <Receiving machine ip>
# Name of the user on the host that connects to the agent.
  username: <Username of receiver>
# Enables the agent to receive packets from the host.
  auth_key: <md5 password>
# The private user password
  priv_key: <des password>

# For v2_endpoint:

# For more hosts you can add more entry with endpoint2, endpoint3, etc

endpoint1:
# Name or IP address of the remote SNMP host.
  host_ip: <Receiving machine ip>
  community: <community name>
```

- Restart the tendrl-notifier service:

```
systemctl restart tendrl-notifier
```