



Red Hat Gluster Storage 3.4

Console Administration Guide

System Administration of Red Hat Gluster Storage Environments using the
Administration Portal

Red Hat Gluster Storage 3.4 Console Administration Guide

System Administration of Red Hat Gluster Storage Environments using the Administration Portal

Red Hat Gluster Storage Documentation Team
Customer Content Services

Legal Notice

Copyright © 2015-2018 Red Hat Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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, 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 Software Collections 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 is a step-by-step guide for users to configure and manage Red Hat Gluster Storage environment using the Administration Portal. This guide is intended for advanced users, and assumes that you have successfully installed the Red Hat Gluster Storage Console and have an understanding of your storage server resources. It describes how to use the Administration Portal, and manage system components and storage infrastructure.

Table of Contents

CHAPTER 1. INTRODUCTION	4
1.1. SYSTEM COMPONENTS	4
1.2. RED HAT GLUSTER STORAGE CONSOLE RESOURCES	4
1.3. ADMINISTRATION OF THE RED HAT GLUSTER STORAGE CONSOLE	5
PART I. THE RED HAT GLUSTER STORAGE CONSOLE INTERFACE	7
CHAPTER 2. GETTING STARTED	8
2.1. GRAPHICAL USER INTERFACE	8
2.2. SEARCH	11
2.3. TAGS	15
CHAPTER 3. DASHBOARD OVERVIEW	17
3.1. VIEWING CLUSTER SUMMARY	17
PART II. MANAGING SYSTEM COMPONENTS	19
CHAPTER 4. MANAGING CLUSTERS	20
4.1. CLUSTER PROPERTIES	20
4.2. CLUSTER OPERATIONS	22
4.3. CLUSTER ENTITIES	25
4.4. CLUSTER PERMISSIONS	27
CHAPTER 5. LOGICAL NETWORKS	29
5.1. INTRODUCTION TO LOGICAL NETWORKS	29
5.2. REQUIRED NETWORKS, OPTIONAL NETWORKS	30
5.3. LOGICAL NETWORK TASKS	30
5.4. LOGICAL NETWORKS AND PERMISSIONS	37
CHAPTER 6. MANAGING RED HAT GLUSTER STORAGE HOSTS	38
6.1. HOSTS PROPERTIES	38
6.2. HOSTS OPERATIONS	38
6.3. MAINTAINING HOSTS	48
6.4. HOSTS ENTITIES	49
6.5. HOSTS PERMISSIONS	51
CHAPTER 7. MANAGING VOLUMES	53
7.1. CREATING A VOLUME	53
7.2. STARTING VOLUMES	55
7.3. CONFIGURING VOLUME OPTIONS	55
7.4. STOPPING VOLUMES	57
7.5. DELETING VOLUMES	58
7.6. MANAGING BRICKS	58
7.7. VOLUMES PERMISSIONS	65
7.8. REBALANCING VOLUME	66
CHAPTER 8. MANAGING GLUSTER HOOKS	69
8.1. VIEWING THE LIST OF HOOKS	69
8.2. VIEWING THE CONTENT OF HOOKS	69
8.3. ENABLING OR DISABLING HOOKS	70
8.4. REFRESHING HOOKS	70
8.5. RESOLVING CONFLICTS	71
CHAPTER 9. MANAGING SNAPSHOTS	75

9.1. CREATING SNAPSHOTS	75
9.2. CONFIGURING SNAPSHOTS	76
9.3. SCHEDULING SNAPSHOTS	77
9.4. RESTORING SNAPSHOTS	80
9.5. ACTIVATING SNAPSHOTS	80
9.6. DEACTIVATING SNAPSHOTS	81
9.7. DELETING SNAPSHOTS	81
CHAPTER 10. MANAGING GEO-REPLICATION	82
10.1. GEO-REPLICATION OPERATIONS	82
CHAPTER 11. USERS	88
11.1. DIRECTORY SERVICES SUPPORT IN RED HAT GLUSTER STORAGE CONSOLE	88
11.2. AUTHORIZATION MODEL	90
11.3. USER PROPERTIES	91
11.4. ASSIGNING AN ADMINISTRATOR OR USER ROLE TO A RESOURCE	93
11.5. REMOVING AN ADMINISTRATOR OR USER ROLE FROM A RESOURCE	94
11.6. USERS OPERATIONS	95
11.7. EVENT NOTIFICATIONS	97
PART III. MONITORING	101
CHAPTER 12. MONITORING RED HAT GLUSTER STORAGE CONSOLE	102
12.1. VIEWING THE EVENT LIST	102
12.2. VIEWING ALERT INFORMATION	103
CHAPTER 13. MONITORING RED HAT GLUSTER STORAGE USING NAGIOS	104
13.1. CONFIGURING NAGIOS	104
13.2. CONFIGURING NAGIOS SERVER FOR SEND MAIL NOTIFICATIONS	106
13.3. VERIFYING THE CONFIGURATION	109
13.4. USING NAGIOS SERVER GUI	110
13.5. MONITORING HOST AND CLUSTER UTILIZATION	128
13.6. TROUBLESHOOTING NAGIOS	131
PART IV. MANAGING ADVANCED FUNCTIONALITY	138
CHAPTER 14. MANAGING MULTILEVEL ADMINISTRATION	139
14.1. CONFIGURING ROLES	139
CHAPTER 15. BACKING UP AND RESTORING THE RED HAT GLUSTER STORAGE CONSOLE	143
15.1. BACKING UP AND RESTORING THE RED HAT GLUSTER STORAGE CONSOLE	143
APPENDIX A. UTILITIES	151
A.1. DOMAIN MANAGEMENT TOOL	151
APPENDIX B. CHANGING PASSWORDS IN RED HAT GLUSTER STORAGE CONSOLE	155
B.1. CHANGING THE PASSWORD FOR THE ADMINISTRATOR USER	155
APPENDIX C. SEARCH PARAMETERS	156
C.1. SEARCH QUERY SYNTAX	156
C.2. SEARCHING FOR RESOURCES	156
C.3. SAVING AND ACCESSING QUERIES AS BOOKMARKS	161
APPENDIX D. CONFIGURATION FILES	163
D.1. NAGIOS CONFIGURATION FILES	163
APPENDIX E. REVISION HISTORY	165

CHAPTER 1. INTRODUCTION

Red Hat Gluster Storage Console is management infrastructure that enables you to create a powerful, scalable storage environment.

It provides IT departments with the tools to meet the challenges of managing complex environments, and enables administrators to reduce the cost and complexity of large deployments. Red Hat Gluster Storage Console includes:

- Support to quickly create and manage Red Hat Gluster Storage trusted storage pool and volumes.
- Multilevel administration to enable administration of physical infrastructure and virtual objects.



NOTE

As of Red Hat Gluster Storage 3.4, the existing Red Hat Gluster Storage Console management infrastructure is supported through the current Red Hat Gluster Storage 3.x life cycle, which ends on October 31, 2019. Red Hat Gluster Storage Web Administration is now the recommended monitoring tool for Red Hat Storage Gluster Clusters. For more information regarding Red Hat Gluster life cycle, see [Red Hat Gluster Storage life Cycle](#)

For Red Hat Gluster Storage Web Administration installation instructions, see the [Red Hat Gluster Storage Web Administration Quick Start Guide](#) and for instructions to monitor your Gluster servers, see the [Red Hat Gluster Storage Web Administration Monitoring Guide](#)

1.1. SYSTEM COMPONENTS

The various components work together seamlessly to enable the system administrator to set up, configure, and maintain the storage environment via an intuitive graphical user interface.

1.1.1. Components

Red Hat Gluster Storage consists of one or more servers and at least one console. The system and all its components are managed through a centralized management system.

1.1.2. The Console

Red Hat Gluster Storage Console is a service that runs on a Red Hat Enterprise Linux 6.5, Red Hat Enterprise Linux 6.6, and Red Hat Enterprise Linux 6.7 servers, providing interfaces for controlling the Red Hat Gluster Storage. It manages user session login and logout, high availability and clustering systems. This is also referred to as the *engine*.

1.1.3. Hosts

Red Hat Gluster Storage Server is a trusted network of storage servers. When you start the first host, the storage pool consists of that host alone. You can add additional storage hosts to the cluster. Red Hat Gluster Storage volumes are created on these clusters. A cluster and all its components can be managed through the Console when a console agent (the VDSM service) is running on each member of the cluster.

1.2. RED HAT GLUSTER STORAGE CONSOLE RESOURCES

The Red Hat Gluster Storage Console manages the following resources within the management infrastructure to create a powerful, scalable storage environment.

- **Hosts** - A host is a physical host (a physical machine) running Red Hat Gluster Storage 3.4. Servers are grouped into storage clusters. Red Hat Gluster Storage volumes are created on these clusters. The system and all its components are managed through a centralized management system.
- **Clusters** - A cluster is a group of linked computers that work together closely, thus in many respects forming a single computer. Hosts in a cluster share the same network infrastructure and the same storage.
- **User** - Red Hat Gluster Storage supports multiple levels of administrators and users with distinct levels of permissions. System administrators can manage and administer objects of the physical infrastructure, such as clusters, hosts, and volume.
- **Events and Monitors** - Alerts, warnings, and other notices about activities within the system help the administrator to monitor the performance and operation of various resources.

1.3. ADMINISTRATION OF THE RED HAT GLUSTER STORAGE CONSOLE

This section provides a high level overview of the tasks and responsibilities of a system administrator for the Red Hat Gluster Storage Console. The tasks are divided into two general groups:

- Configuring a new logical cluster is the most important task of the system administrator. Designing a new cluster requires an understanding of capacity planning and definition of requirements. This is typically determined by the solution architect, who provides the requirements to the system architect. Preparing to set up the storage environment is a significant part of the setup, and is usually part of the system administrator's role.
- Maintaining the cluster, including performing updates and monitoring usage and performance to keep the cluster responsive to changing needs and loads.

The procedures to complete these tasks are described in detail in later sections of this guide.

It is assumed that you have already read the material in *Red Hat Gluster Storage Console 3.4 Installation Guide*.

1.3.1. Maintaining the Red Hat Gluster Storage Console

This section describes how to maintain a Red Hat Gluster Storage Console.

The administrator's tasks include:

- Managing hosts and other physical resources.
- Managing the storage environment. This includes creating, deleting, expanding and shrinking volumes and clusters.
- Monitoring overall system resources for potential problems such as an extreme load on one of the hosts, insufficient memory or disk space, and taking any necessary actions.
- Managing user setup and access, and setting user and administrator permission levels. This includes assigning or customizing roles to suit the needs of the enterprise.

- Troubleshooting for specific users or hosts or for overall system functionality.

These tasks are described in detail in later sections of this guide.

PART I. THE RED HAT GLUSTER STORAGE CONSOLE INTERFACE

CHAPTER 2. GETTING STARTED

The Administration Portal allows you to create, monitor and maintain your Red Hat Gluster Storage using an interactive graphical user interface (GUI). The GUI functions in two modes - tree, or flat - allowing you to navigate the system's resources either hierarchically or directly. The powerful search feature enables you to locate any resource in the enterprise, wherever it may be in the hierarchy, and you can use tags and bookmarks to help you to store the results of your searches for later reference.

It is assumed that you have correctly installed Red Hat Gluster Storage Console, including hosts, and have logged into the Administration Portal. If you are attempting to set up Red Hat Gluster Storage Console, see *Red Hat Gluster Storage Console 3.4 Installation Guide*

2.1. GRAPHICAL USER INTERFACE

After you have successfully logged into Red Hat Gluster Storage Console, the Administration Portal displays. The GUI consists of a number of contextual panes and menus, and can be used in two modes - tree mode, and flat mode. Tree mode allows you to browse the object hierarchy of a cluster, and is the recommended manner of operation. The elements of the GUI are shown in the figure below.




Figure 2.1. Graphical User Interface Elements of the Administration Portal

Graphical User Interface Elements

- 1 **Header**

The **Header** bar contains the name of the current logged-in user, the **Sign Out** button, the **About** button, and the **Configure** button. The **About** button provides access to version information. The **Configure** button allows you to configure user roles.

- 2 **Search Bar**

The **Search** bar allows you to quickly search for resources such as hosts and volumes. You can build queries to find the resources that you need. Queries can be as simple as a list of all the hosts in the system, or much more complex. As you type each part of the search query, you will be offered choices to assist you in building the search. The star icon can be used to save the search as a bookmark.

- **3 Resource Tabs**

All resources, such as hosts and clusters, can be managed using the appropriate tab. Additionally, the **Events** tab allows you to manage and view events across the entire system.

Clicking a tab displays the results of the most recent search query on the selected object. For example, if you recently searched for all hosts starting with "M", clicking the **Hosts** tab displays a list of all hosts starting with "M".

The Administration Portal provides the following tabs: **Clusters**, **Hosts**, **Volumes**, **Users**, and **Events**.

- **4 Results List**

Perform a task on an individual item, multiple items, or all the items in the results list, by selecting the items and then clicking the relevant action button. If multiple selection is not possible, the button is disabled.

Details of a selected item display in the details pane.

- **5 Details Pane**


The **Details** pane displays detailed information about a selected item in the Results Grid. If multiple items are selected, the **Details** pane displays information on the first selected item only.

- **6 Bookmarks Pane**

Bookmarks are used to save frequently used or complicated searches for repeated use. Bookmarks can be added, edited, or removed.

- **7 Alerts/Events Pane**

The **Alerts** pane lists all events with a severity of **Error** or **Warning**. The system records all events, which are listed as audits in the **Alerts** section. Like events, alerts can also be viewed in the lowermost panel of the **Events** tab by resizing the panel and clicking the **Alerts** tab. This tabbed panel also appears in other tabs, such as the **Hosts** tab.



IMPORTANT

The minimum supported resolution for viewing the Administration Portal in a web browser is 1024 x 768. When viewed at a lower resolution, the Administration Portal will not render correctly.

2.1.1. Tree Mode and Flat Mode

The Administration Portal provides two different modes for managing your resources - tree mode, and flat mode.

Tree mode displays resources in a hierarchical view for each cluster, from the highest level of the cluster down to the individual volumes. Tree mode provides a visual representation of the storage system. Working in tree mode is recommended for most operations.

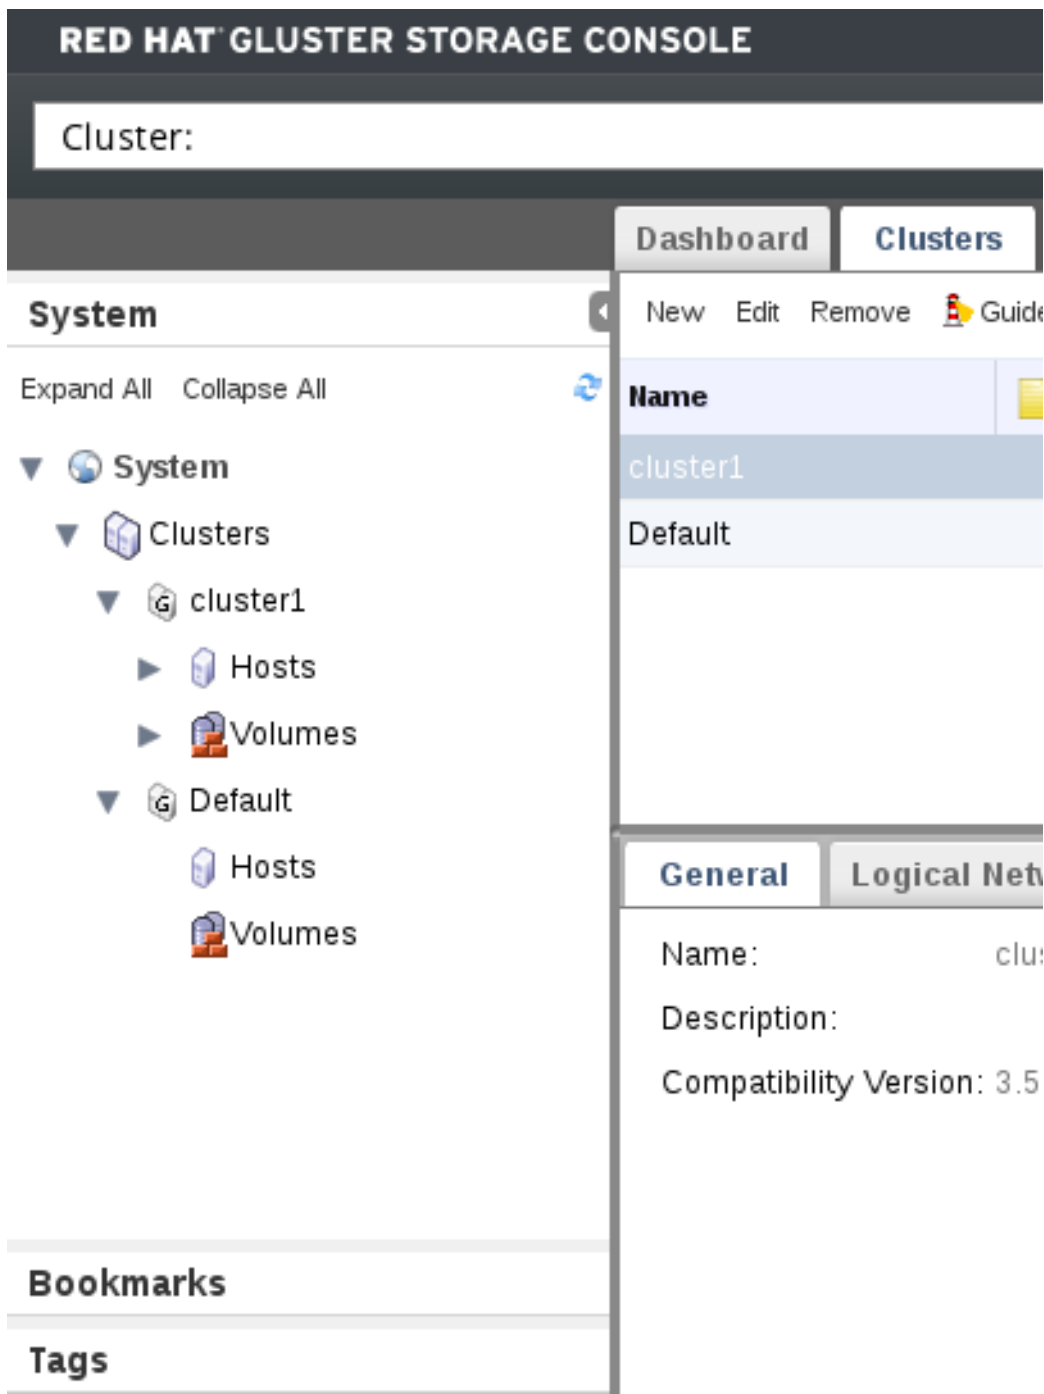


Figure 2.2. Tree Mode

Flat mode offers powerful search functionality, and allows you to customize how you manage your system. It gives you access to any resource, regardless of its position in the enterprise. In this mode, the full power of the search feature can be used. Flat mode does not limit you to viewing the resources of a single hierarchy, allowing you to search across clusters. For example, flat mode makes it possible to find all hosts that are using more than 80% CPU across clusters, or locate all hosts that have the highest utilization. In addition, certain objects are not in the cluster hierarchy, so they will not appear in tree mode. For example, users are not part of the cluster hierarchy, and can be accessed only in flat mode.

To access flat mode, click **System** in the pane on the left-hand side of the screen.

Name	Compatibility Version	Description	Host Count	VM Count
cluster1	3.5		2	0
Default	3.5	The default server cluster	0	0

Name:	cluster1	Total No. Of Volumes:	0
Description:		No. Of Volumes Up:	0
Compatibility Version:	3.5	No. Of Volumes Down:	0

Figure 2.3. Flat Mode

2.2. SEARCH

The Administration Portal is designed to enable the management of thousands of resources, such as volumes, hosts, and clusters. When managing the storage environment, it is recommended that large lists of resources, such as volumes, are reduced to a manageable number (for example, 10). This allows tasks to be performed on a smaller list, or to select specific items on the list on which to perform a given task.

To perform a search, enter the search query (free-text or syntax-based) in the **Search** bar at the top of the Administration Portal. Search queries can be saved as bookmarks for future reuse ([Section C.3, “Saving and Accessing Queries as Bookmarks”](#)). This eliminates the need to re-enter a search query each time specific search results are required.

2.2.1. Search Syntax

The syntax of search queries for Red Hat Gluster Storage Console resources is as follows:

```
result-type: {criteria} [sortby sort_spec]
```

Syntax Examples

The following examples describe how search queries are used, and help you to understand how Red Hat Gluster Storage Console assists with building search queries.

Table 2.1. Example Search Queries

Example	Result
Volumes: status = up	Displays a list of all volumes that are up.
Volumes: cluster = data	Displays a list of all volumes of the cluster data.
Events: severity > normal sortby time	Displays the list of all events whose severity is higher than Normal , sorted by time.

2.2.1.1. Auto-Completion

The Administration Portal provides auto-completion to help you create valid and powerful search queries. As you type each part of a search query, a drop-down list of choices for the next part of the search opens below the **Search** bar. You can select from the list and then continue typing or selecting the next part of the search, or ignore the options and continue entering your query manually.

The following table provides examples of how Administration Portal auto-completion assists in constructing a query:

Volumes: status = down

Table 2.2. Example Search Queries using Auto-Completion

Input	List Items Displayed	Action
v	Volumes (1 option only)	Select Volumes or; Type Volumes
Volumes:	All volumes properties	Type s
Volumes: s	volume properties starting with s	Select status or type status
Volumes: status	= !=	Select or type =
Volumes: status =	All status values	Select or type down

2.2.1.2. Result-Type Options

The result type allows you to search for resources of any of the following types:

- **Host** for a list of hosts
- **Event** for a list of events
- **Users** for a list of users
- **Cluster** for a list of clusters
- **Volumes** for a list of volumes

Each type of resource has a unique set of properties and a set of other resource types that it is associated with, so each search type has a set of valid syntax combinations. However, using the auto-complete feature helps you to easily create valid queries.

2.2.1.3. Search Criteria

You can specify the search criteria after the colon in the query. The syntax of **{criteria}** is as follows:

<prop> <operator> <value>

or

`<obj-type>.<prop> <operator> <value>`

Examples

The following table describes the parts of the syntax:

Table 2.3. Example Search Criteria

Part	Description	Values	Example	Note
prop	The property of the searched-for resource. Can also be the property of a resource type (see obj - type), or <i>tag</i> (custom tag).	See the information for each of the search types in Section 2.2.1.3.1, “Wildcards and Multiple Criteria” .	Status	--
obj-type	A resource type that can be associated with the searched-for resource.	See the explanation of each of the search types in Section 2.2.1.3.1, “Wildcards and Multiple Criteria” .	Users	--
operator	Comparison operators.	= != (not equal) > < >= <=	--	Value options depend on obj-type.

Part	Description	Values	Example	Note
Value	What the expression is being compared to.	String Integer Ranking Date (formatted according to regional settings)	Jones 256 normal	<ul style="list-style-type: none"> • Wildcards can be used within strings. • "" (two sets of quotation marks with no space between them) can be used to represent an uninitialized (empty) string. • Double quotation marks should be used around a string or date that contains spaces.

2.2.1.3.1. Wildcards and Multiple Criteria

Wildcards can be used in the `<value>` part of the syntax for strings. For example, to find all users beginning with `m`, enter `m*`.

You can perform a search with two criteria by using the Boolean operators **AND** and **OR**. For example:

Volumes: name = m* AND status = Up

This query returns all volumes whose names begin with "m".

When two criteria are specified without **AND** or **OR**, **AND** is implied. **AND** precedes **OR**, and **OR** precedes implied **AND**.

2.2.1.4. Determining Sort Order

You can determine the sort order of the returned information by using **sortby**. Sort direction (**asc** for ascending, **desc** for descending) can be included.

For example:

```
events: severity > normal sortby time desc
```

This query returns all events whose severity is higher than Normal, sorted by time (descending order).

2.2.2. Saving Queries as Bookmarks

Search queries can be saved as bookmarks. This allows you to sort and display results lists with a single click. You can save, edit and remove bookmarks with the Bookmarks pane. See [Section C.3, “Saving and Accessing Queries as Bookmarks”](#) for details.

2.3. TAGS

After your Red Hat Gluster Storage is set up and configured to your requirements, you can customize the way you work with it using tags. Tags provide one key advantage to system administrators - they allow system resources to be arranged into groups or categories. This is useful when many objects exist in the storage environment and the administrator would like to concentrate on a specific set of them.

This section describes how to create and edit tags, assign them to hosts and search using the tags as criteria. Tags can be arranged in a hierarchy that matches a structure, to fit the requirements of the enterprise.

Administration Portal tags can be created, modified, and removed using the **Tags** pane.

Procedure 2.1. Creating a tag

1. In tree mode or flat mode, click the resource tab for which you wish to create a tag. For example, **Hosts**.
2. Click the **Tags** tab. Select the node under which you wish to create the tag. For example, click the root node to create it at the highest level. The **New** button is enabled.
3. Click **New** at the top of the **Tags** pane. The **New Tag** dialog box displays.
4. Enter the **Name** and **Description** of the new tag.
5. Click **OK**. The new tag is created and displays on the **Tags** tab.

Procedure 2.2. Modifying a tag

1. Click the **Tags** tab. Select the tag that you wish to modify. The buttons on the **Tags** tab are enabled.
2. Click **Edit** on the **Tags** pane. The **Edit Tag** dialog box displays.
3. You can change the **Name** and **Description** of the tag.
4. Click **OK**. The changes in the tag display on the **Tags** tab.

Procedure 2.3. Deleting a tag

1. Click the **Tags** tab. The list of tags will display.
2. Select the tags to be deleted and click **Remove**. The **Remove Tag(s)** dialog box displays.

3. The tags are displayed in the dialog box. Check that you are sure about the removal. The message warns you that removing the tags will also remove all descendants of the tags.
4. Click **OK**. The tags are removed and no longer display on the **Tags** tab. The tags are also removed from all the objects to which they were attached.

Tags can be attached to hosts and users.

Procedure 2.4. Adding or removing a tag to or from one or more object instances

1. Search for the objects that you wish to tag or untag so that they are among the objects displayed in the results list.
2. Select one or more objects on the results list.
3. Click the **Assign Tags** button on the tool bar or right-click menu option.
4. A dialog box provides a list of tags. Select the check box to assign a tag to the object, or deselect the check box to detach the tag from the object.
5. Click **OK**. The specified tag is now added or removed as a custom property of the selected objects.

A user-defined tag can be a property of any object (for example, a host), and a search can be conducted to find it.

To search for objects using tags:

- Follow the search instructions in [Section 2.2, "Search"](#) , and enter a search query using "tag" as the property and the desired value or set of values as criteria for the search.

The objects tagged with the tag criteria that you specified are listed in the results list.

CHAPTER 3. DASHBOARD OVERVIEW

Red Hat Gluster Storage Console Dashboard displays an overview of all the entities in Red Hat Gluster Storage like Hosts, Volumes, Bricks, and Clusters. The Dashboard shows a consolidated view of the system and helps the administrator to know the status of the system. Listed below are the Dashboard items in Red Hat Gluster Storage Console:

- **Capacity:** Displays the Total, Used and Available storage capacity in the system. Its calculated by aggregating data from all the hosts in the System.
- **Utilization:** Displays the average usage percentage of CPU, Memory. This is averaged across all the hosts in system.
- **Alerts :** Displays the number of alerts in the system. The Alerts tab will display a red exclamation icon if there is an alert or alerts. Click on the alerts arrow icon to open the alerts dialog box. To delete the alert/alerts, click the cross icon at the right hand side of the dialogue box.
- **Hosts:** Displays the total number of Hosts in the system and the number of hosts in down state.
- **Volumes:** Displays the total number of volumes in the system across all cluster and number of them in UP, DOWN, Degraded, Partial, or Stopped status.
- **NICs:** Displays the number of network interfaces in the hosts.
- **Network:** Displays the transmission and receiving rate of the NICs.



NOTE

Top Utilization popup: To see the top utilizers for Capacity, Utilization, and Network, hover over the individual tabs and a popup will appear with the top utilization data.

3.1. VIEWING CLUSTER SUMMARY

Procedure 3.1. Viewing Cluster Summary

1. In **Dashboard** tab, select the cluster name from drop-down list to view cluster capacity details of a specific cluster.

Select **All Clusters** to view cluster capacity details of all clusters.

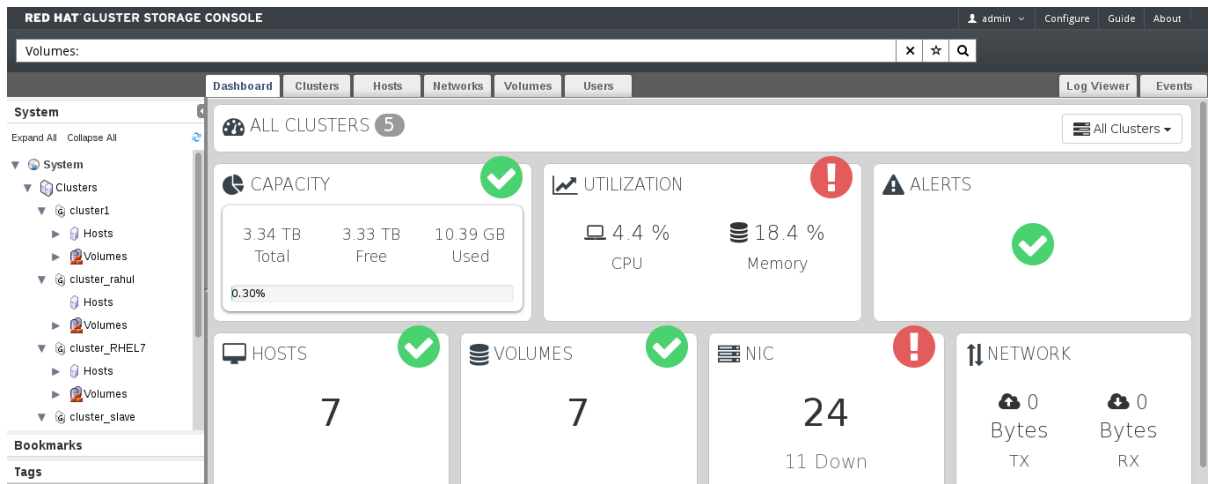


Figure 3.1. Dashboard Overview

2. View the cluster and volume details by hovering over each dashboard item.
3. Click **CAPACITY** or **VOLUMES** to view the cluster and volume details respectively.
4. Click **UTILIZATION**, **HOST**, **NIC**, or **NETWORK** to view the host and network details.



NOTE

Clicking each item takes you to corresponding tab in Red Hat Gluster Storage Console.

PART II. MANAGING SYSTEM COMPONENTS

CHAPTER 4. MANAGING CLUSTERS

The cluster is the highest level entity for all physical and logical resources within a storage environment. This chapter describes how to create and manage clusters.

4.1. CLUSTER PROPERTIES

Use the **Clusters** tab in the Administration Portal to define, manage, and view clusters.

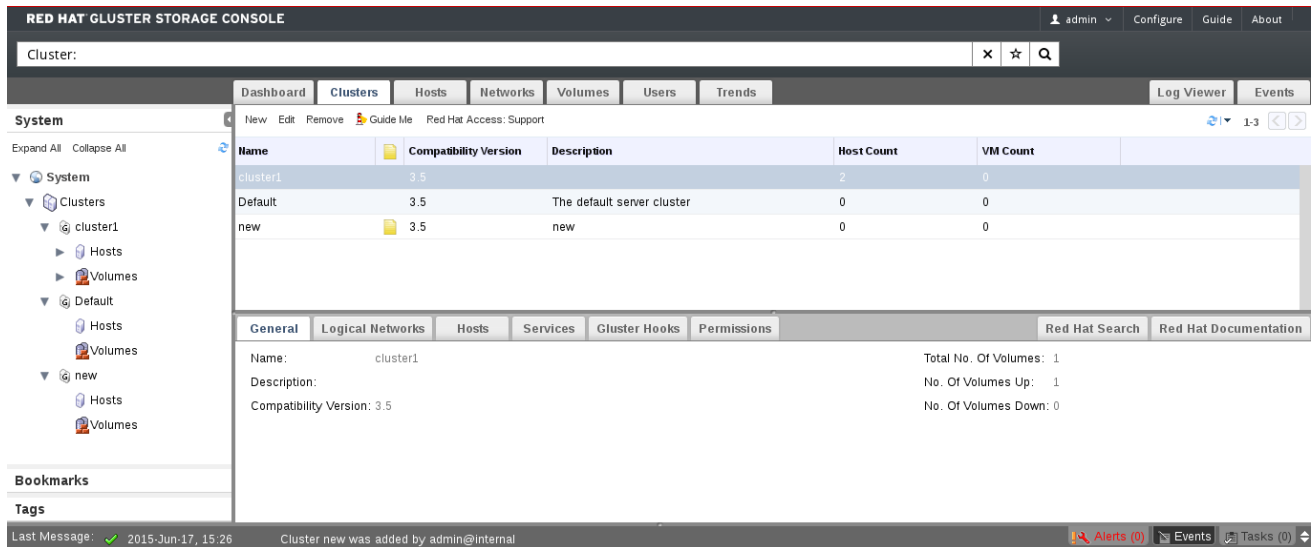


Figure 4.1. Clusters Tab

The following table describes the cluster properties displayed in the **New Cluster** and **Edit Cluster** dialog boxes. Missing mandatory fields and invalid entries are outlined in red when you click **OK** to close the **New Cluster** or **Edit Cluster** dialog box.

Table 4.1. Cluster Properties

Field	Description
Name	The name of the cluster. This must be a unique name and may use any combination of uppercase or lowercase letters, numbers, hyphens and underscores. Maximum length is 40 characters. The name can start with a number and this field is mandatory.
Description	The description of the cluster. This field is optional, but recommended.


Field	Description
Compatibility Version	<p>The version of Red Hat Gluster Storage Console with which the cluster is compatible. All hosts in the cluster must support the indicated version.</p> <ul style="list-style-type: none"> • Clusters with compatibility version 3.2 can manage Red Hat Gluster Storage 2.1 nodes. • Clusters with compatibility version 3.4 can manage Red Hat Gluster Storage 2.1 Update 2 nodes. • Clusters with compatibility version 3.4 can manage Red Hat Gluster Storage 3.0 nodes. <p> NOTE</p> <p>The default compatibility version is 3.4.</p>

Table 4.2. Compatibility Matrix

Feature	Compatibility Version 3.2	Compatibility Version 3.3	Compatibility Version 3.4
View advanced details of a particular brick of the volume through the Red Hat Gluster Storage Console.	Supported	Supported	Supported
Synchronize brick status with the engine database.	Supported	Supported	Supported
Manage glusterFS hooks through the Red Hat Gluster Storage Console. View the list of hooks available in the hosts, view the contents and status of hooks, enable or disable hooks, and resolve hook conflicts.	Supported	Supported	Supported
Display Services tab with NFS and SHD service status.	Supported	Supported	Supported
Manage volume rebalance through the Red Hat Gluster Storage Console. Rebalance volume, stop rebalance, and view rebalance status.	Not Supported	Supported	Supported
Manage remove-brick operations through the Red Hat Gluster Storage Console. Remove-brick, stop remove-brick, view remove-brick status, and retain the brick being removed.	Not Supported	Supported	Supported

Feature	Compatibility Version 3.2	Compatibility Version 3.3	Compatibility Version 3.4
Allow using system's root partition for bricks and re-using the bricks by clearing the extended attributes.	Not Supported	Supported	Supported
Addition of RHS U2 nodes	Not Supported	Supported	Supported
Viewing Nagios Monitoring Trends	Not Supported	Not Supported	Supported

4.2. CLUSTER OPERATIONS

4.2.1. Creating a New Cluster

You can create a new cluster using the **New** option in the **Clusters** tab.

Procedure 4.1. To Create a New Cluster

1. Open the **Clusters** view by expanding the **System** tab and selecting the **Cluster** tab in the **Tree** pane. Alternatively, select **Clusters** from the **Details** pane.
2. Click **New** to open the **New Cluster** dialog box.

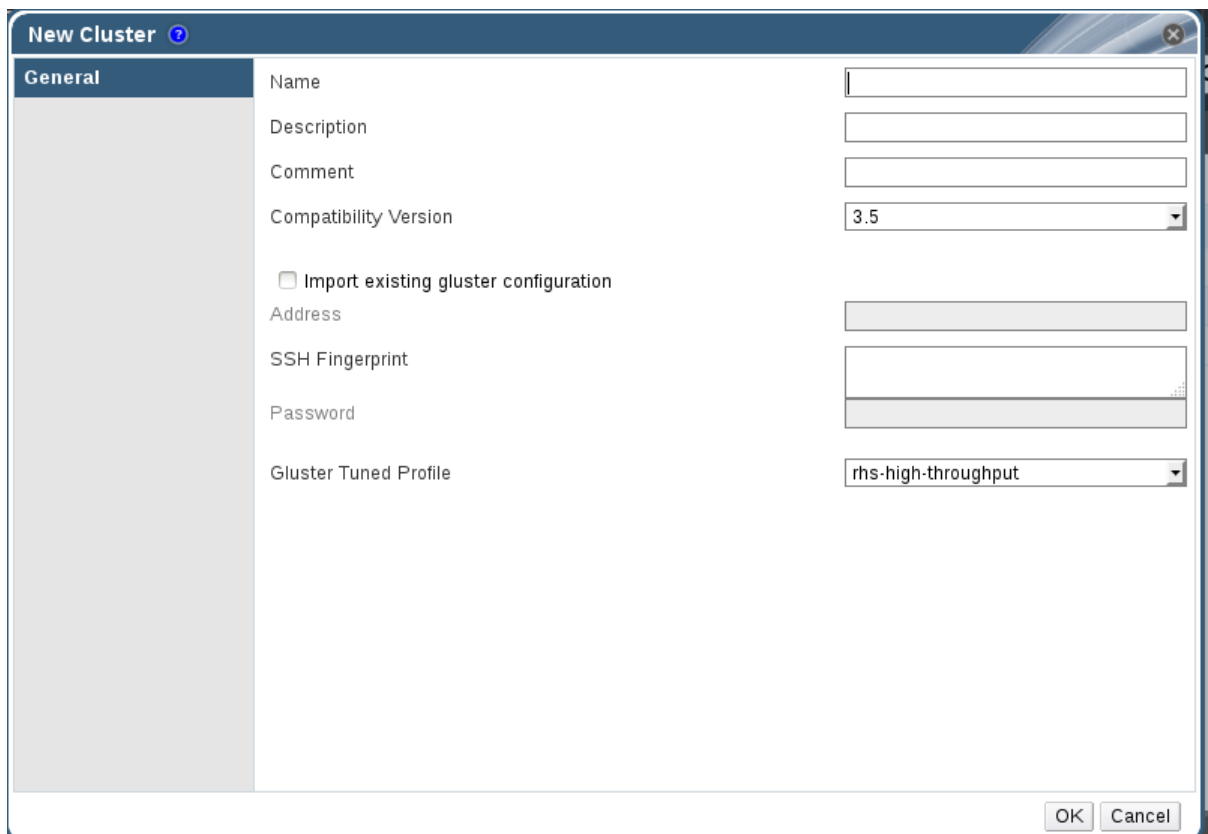


Figure 4.2. New Cluster Dialog Box

3. Enter the cluster **Name**, **Description** and **Compatibility Version**. The name cannot

include spaces. When the user selects **Import existing gluster configuration** and enters the **Address**, fingerprint will be fetched automatically by the Red Hat Gluster Storage Console.

4. Click **OK** to create the cluster. The new cluster displays in the **Clusters** tab.
5. Click **Guide Me** to configure the cluster. The **Guide Me** window lists the entities you need to configure for the cluster. Configure these entities or postpone configuration by clicking **Configure Later**. You can resume the configuration process by selecting the cluster and clicking **Guide Me**. To import an existing cluster, see [Section 4.2.2, “Importing an Existing Cluster”](#).

Gluster Tuned Profile

Tuned profiles help to enhance the performance of the system by applying some predefined set of system parameters. We have two tuned profiles for Red Hat Gluster Storage:

- **rhs-high-throughput**: It is a default profile which will be applied on the RHGS nodes. This helps to enhance the performance of RHGS Volume.
- **rhs-virtualization**: If the number of clients is greater than 100, you must switch to the rhs-virtualization tuned profile. For more information, see *Number of Clients* in *Red Hat Gluster Storage Administration Guide*.

4.2.2. Importing an Existing Cluster

You can import a Red Hat Gluster Storage cluster and all the hosts belonging to the cluster into the Red Hat Gluster Storage Console.

When you provide details such as the IP address or host name and password of any host in the cluster, the gluster **peer status** command executes on that host through SSH, then displays a list of hosts that are part of the cluster. You must manually verify the fingerprint of each host and provide passwords for them. If some hosts are not reachable, then import cluster will not add these hosts to the cluster during import.

Procedure 4.2. To Import an Existing Cluster

1. In the **Tree** pane, click **System** tab, then click the **Clusters** tab.
2. Click **New** to open the **New Cluster** dialog box.
3. Enter the cluster **Name**, **Description** and **Compatibility Version**. The name cannot include spaces.
4. Select **Import existing gluster configuration** to import the cluster.
5. In the **Address** field, enter the host name or IP address of a host in the cluster.

The host **Fingerprint** displays to indicate the connection host. If a host is unreachable or if there is a network error, **Error in fetching fingerprint** displays in the **Fingerprint** field.

6. Enter the **Root Password** for the host in the **Password** field and click **OK**.
7. The **Add Hosts** window opens, and a list of hosts that are part of the cluster displays.

8. For each host, enter the **Name** and **Root Password**. If you wish to use the same password for all hosts, check **Use a common password** and enter a password.
9. Click **Apply** to set the password for all hosts then click **OK** to submit the changes.

4.2.3. Editing a Cluster

Procedure 4.3. To Edit a Cluster

1. Click the **Clusters** tab to display the list of host clusters. Select the cluster that you want to edit.
2. Click **Edit** to open the **Edit Cluster** dialog box.
3. Enter a **Name** and **Description** for the cluster and select the compatibility version from the **Compatibility Version** drop down list.
4. Click **OK** to confirm the changes and display the host cluster details.

4.2.4. Viewing Hosts in a Cluster

Procedure 4.4. To View Hosts in a Cluster

1. Click the **Clusters** tab to display a list of host clusters. Select the desired cluster to display the **Details** pane.
2. Click the **Hosts** tab to display a list of hosts.

The screenshot shows the Red Hat Cluster Storage Console interface. The main window displays the 'Hosts' tab for a cluster named 'new'. The interface includes a navigation pane on the left, a main content area with a table of hosts, and a bottom status bar.

Name	Compatibility Version	Description	Host Count	VM Count
cluster1	3.5		2	0
Default	3.5	The default server cluster	0	0
new	3.5	new	0	0

Name	Hostname:IP	Status
↑ rhs-client24.lab.eng.blr.redhat.com	rhs-client24.lab.eng.blr.redhat.com	Up
↑ rhs-client33.lab.eng.blr.redhat.com	rhs-client33.lab.eng.blr.redhat.com	Up

The status bar at the bottom shows: Last Message: 2015-Jun-17, 15:26 Cluster new was added by admin@internal. Alerts: 0, Events, Tasks: 0.

Figure 4.3. The Hosts tab on the Cluster Details pane

4.2.5. Removing a Cluster

You can permanently remove clusters that are not in use. Deleting unused clusters saves system resources, as existing hosts are contacted at regular intervals.

**WARNING**

Red Hat recommends that you do not remove the default cluster.

Procedure 4.5. To Remove a Cluster

1. Click the **Clusters** tab to display a list of clusters. If the required cluster is not visible, perform a search.
2. Select the cluster to be removed. Ensure that there are no running hosts or volumes.
3. Click the **Remove** button.
4. A dialog box lists all the clusters selected for removal. Click **OK** to confirm the removal.

4.3. CLUSTER ENTITIES

You can configure cluster entities using the Console.

Cluster Entities

A cluster is a collection of hosts. The **Hosts** tab displays all information related to the hosts in a cluster.

Table 4.3. Host Tab Properties

Field	Description
Name	The name of the host.
Hostname/IP	The name of the host/IP address.
Status	The status of the host.

Cluster Logical Networks Entities

Logical networks enable hosts to communicate with other hosts, and for the Console to communicate with cluster entities. You must define logical networks for each cluster.

Table 4.4. Cluster Logical Networks Tab Properties

Field	Description
Name	The name of the logical networks in a cluster.
Status	The status of the logical networks.
Role	The hierarchical permissions available to the logical network.

Field	Description
Description	The description of the logical networks.

Cluster Permissions Entities

Cluster permissions define which users and roles can work in a cluster, and what operations the users and roles can perform.

Table 4.5. Cluster Permissions Tab Properties

Field	Description
User	The user name of an existing user in the directory services.
Role	The role of the user. The role comprises of user, permission level and object. Roles can be default or customized roles.
Inherited Permissions	The hierarchical permissions available to the user.

Gluster Hooks

Gluster Hooks are volume lifecycle extensions. You can manage the Gluster Hooks from Red Hat Gluster Storage Console.

Table 4.6. Gluster Hooks Tab Properties

Field	Description
Name	The name of the hook.
Volume Event	Events are instances in the execution of volume commands like create, start, stop, add-brick, remove-brick, set and so on. Each of the volume commands have two instances during their execution, namely Pre and Post. Pre and Post refers to the time just before and after the corresponding volume command has taken effect on a peer respectively.
Stage	When the event should be executed. For example, if the event is <i>Start Volume</i> and the <i>Stage</i> is <i>Post</i> , the hook will be executed after the start of the volume.
Status	Status of the gluster hook.
Content Type	Content type of the gluster hook.

Services

The service running on a host can be searched using the *Services* tab.

Table 4.7. Services Tab Properties


Field	Description
Host	The ip of the host.
Service	The name of the service.
Port	The port number of the host.
Status	The status of the host.
Process Id	The process id of the host.

4.4. CLUSTER PERMISSIONS

A cluster administrator has system administrator permissions for a specific cluster only. This is a hierarchical model, which means that a user assigned the cluster administrator role for a cluster can manage all objects in that cluster. The cluster administrator role permits the following actions:

- Creation and removal of specific clusters.
- Addition and removal of hosts.
- Permission to attach users to hosts within a single cluster.

This is useful when there are multiple clusters, each of which require their own system administrators. A cluster administrator has permissions for the assigned cluster only, not for all clusters.



NOTE

You can only assign roles and permissions to existing users.

Procedure 4.6. To Add a Cluster Administrator Role

1. Click the **Clusters** tab to display the list of clusters. If the required cluster is not visible, perform a search.
2. Select the cluster that you want to edit. Click the **Permissions** tab in the **Details** pane to display a list of existing users and their current roles and inherited permissions.
3. Click **Add** to display the **Add Permission to User** dialog box. Enter all or part of a name or user name in the **Search** box, then click **Go**. A list of possible matches displays in the results list.
4. Select the user you want to modify. Scroll through the **Role to Assign** list and select **ClusterAdmin**.
5. Click **OK** to display the name of the user and their assigned role in the **Permissions** tab.

Procedure 4.7. To Remove a Cluster Administrator Role

1. Click the **Clusters** tab to display a list of clusters. If the required cluster is not visible, perform a search.
2. Select the cluster that you want to edit. Click the **Permissions** tab in the **Details** pane to display a list of existing users and their current roles and inherited permissions.
3. Select the user you want to modify and click **Remove**. This removes the user from the **Permissions** tab and from associated hosts and volumes.

CHAPTER 5. LOGICAL NETWORKS

5.1. INTRODUCTION TO LOGICAL NETWORKS

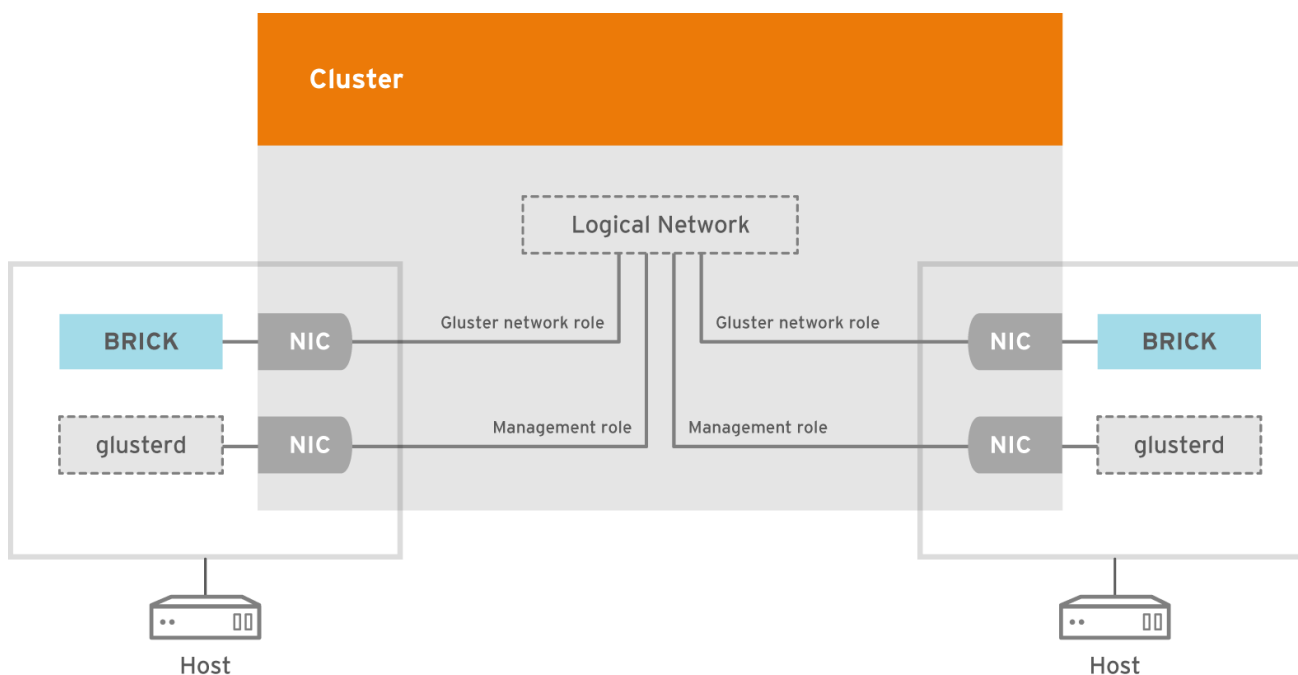
A logical network is a named set of global network connectivity properties in your system. When a logical network is added to a host, it may be further configured with host-specific network parameters. Logical networks optimize network flow by grouping network traffic by usage, type, and requirements.

Logical networks allow both connectivity and segregation. You can create a logical network for gluster storage communication to optimize network traffic between hosts and gluster bricks, a logical network specifically for all network traffic, or multiple logical networks to carry the traffic of groups of networks.

The default logical network is the management network called **ovirtmgmt**. The **ovirtmgmt** network carries all traffic, until another logical network is created. It is meant especially for management communication between the Red Hat Gluster Storage Console and hosts.

When a host has multiple interfaces, user can choose the interface to be used while adding the brick, by tagging one of the host's interface with the Gluster Network role. A logical network that has been designated as **Required** must be configured in all of a cluster's hosts before it is operational.

Optional networks can be used by any host they have been added to.



GLUSTER_334434_0615

Figure 5.1. Logical Network architecture



WARNING

Do not change networking in a cluster if any hosts are running as this risks making the host unreachable.

5.2. REQUIRED NETWORKS, OPTIONAL NETWORKS

Red Hat Gluster Storage Console distinguishes between required networks and optional networks.

Required networks must be applied to all hosts in a cluster for the cluster and network to be **Operational**. By default, Logical networks are added to clusters as **Required** networks.

Click **Manage Networks** to change a network's **Required** designation. Optional networks are those logical networks that have not been explicitly declared **Required** networks. Optional networks can be implemented on only the hosts that use them. The presence or absence of these networks does not affect the **Operational** status of a host.

5.3. LOGICAL NETWORK TASKS

5.3.1. Using the Networks Tab

The **Networks** resource tab provides a central location for users to perform network-related operations and search for networks based on each network's property or association with other resources.

All networks in the Red Hat Gluster Storage environment display in the results list of the **Networks** tab. The **New**, **Edit** and **Remove** buttons allow you to create, change the properties of, and delete logical networks within the system.

Click on each network name and click **Clusters**, **Hosts**, and **Permissions** tabs in the details pane to perform functions including:

- Attaching or detaching the networks to clusters and hosts
- Adding and removing permissions for users to access and manage networks

These functions are also accessible through each individual resource tab.

5.3.2. Creating a New Logical Network in Cluster

Summary

Create a logical network and define its use in a cluster.

Procedure 5.1. Creating a New Logical Network in a Cluster

1. Click the **Networks** or **Clusters** tab in tree mode and select a network or cluster.
2. Click the **Logical Networks** tab of the details pane to list the existing logical networks.
3.
 - From the **Clusters** tab, select **Logical Networks** sub-tab and click **Add Network** to open the **New Logical Network** window.
 - From the **Networks** tab, click **New** to open the **New Logical Network** window.

New Logical Network

The Network will be added to the Data Center Default as well.


Name

Description

Comment

Network Parameters

Network Label

VM network 

MTU

Default (1500)

Custom

OK Cancel

Figure 5.2. New Logical Network

4. Enter a **Name**, **Description**, and **Comment** for the logical network.
5. From the **Cluster** tab, select the clusters to which the network will be assigned. You can also specify whether the logical network will be a required network.

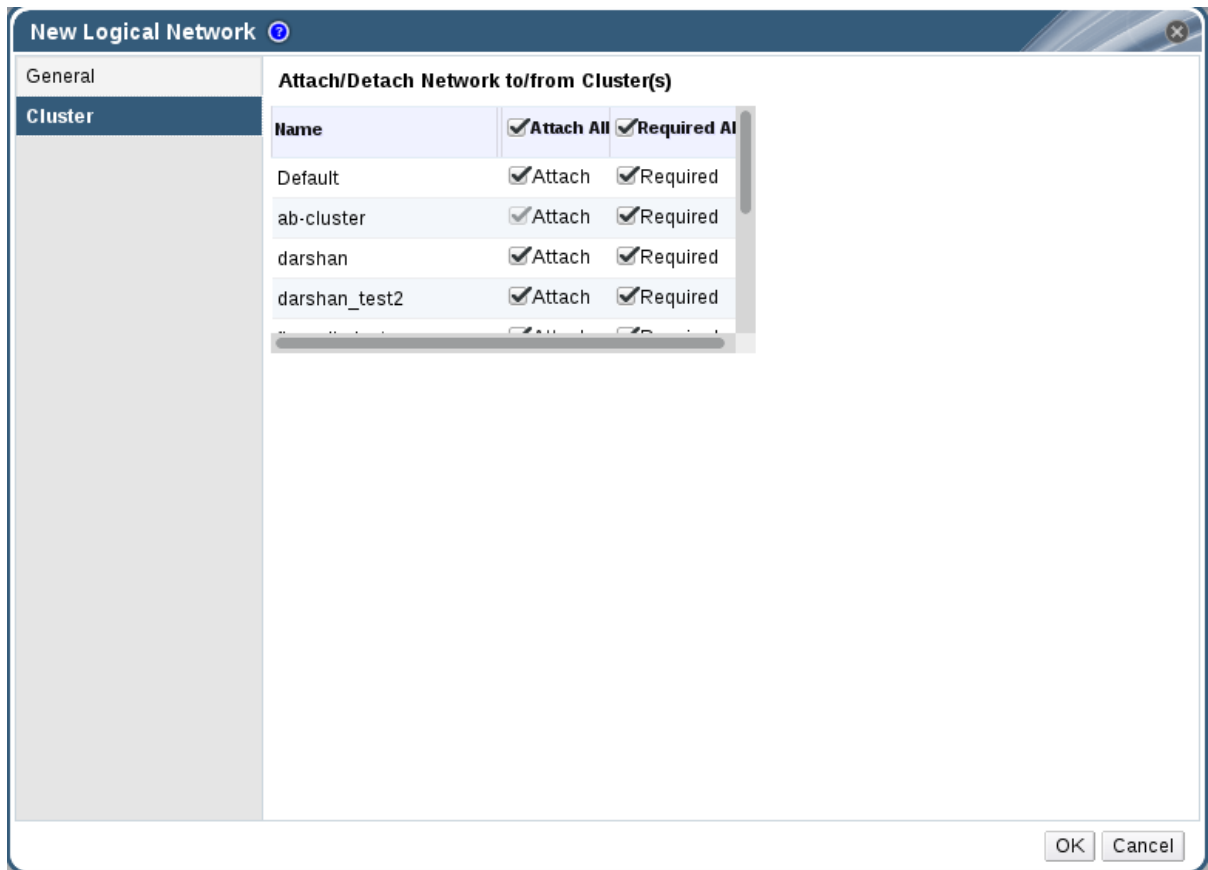


Figure 5.3. New Logical Network - Cluster

6. Click **OK**.

Result

You have defined a logical network as a resource required by a cluster or clusters in the Network. If you entered a label for the logical network, it will be automatically added to all host network interfaces with that label.

5.3.3. Editing a Logical Network

Summary

Edit the settings of a logical network.

Procedure 5.2. Editing a Logical Network

1. Click **Networks** tab in tree mode and select a Network.
2. Select a logical network and click **Edit** to open the **Edit Logical Network** window.

Edit Logical Network ?

General

Data Center: Default

Name: scName

Description: Name

Comment: Name

Edit Network Parameters

Network Label: VM

VM network vm

MTU

Default (1500)

Custom

OK Cancel

Figure 5.4. New Logical Network - Cluster

3. Edit the necessary settings.
4. Click **OK** to save the changes.

Result

You have updated the settings of your logical network.



NOTE

You cannot rename a logical network that is already configured on a host.

5.3.4. Explanation of Settings and Controls in the New Logical Network and Edit Logical Network Windows

5.3.4.1. Logical Network General Settings Explained

The table below describes the settings for the **General** tab of the **New Logical Network** and **Edit Logical Network** window.

Table 5.1. New Logical Network and Edit Logical Network Settings

Field Name	Description
------------	-------------

Field Name	Description
Name	The name of the logical network. This text field has a 15-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
Description	The description of the logical network. This text field has a 40-character limit.
Comment	A field for adding plain text, human-readable comments regarding the logical network.
Network Label	Allows you to specify a new label for the network or select from existing labels already attached to host network interfaces. If you select an existing label, the logical network will be automatically assigned to all host network interfaces with that label.

5.3.4.2. Logical Network Cluster Settings Explained

The table below describes the settings for the **Cluster** tab of the **New Logical Network** window.

Table 5.2. New Logical Network Settings

Field Name	Description
Attach/Detach Network to/from Cluster(s)	<p>Allows you to attach or detach the logical network from clusters and specify whether the logical network will be a required network for individual clusters.</p> <p>Name - the name of the cluster to which the settings will apply. This value cannot be edited.</p> <p>Attach All - Allows you to attach or detach the logical network to or from all clusters. Alternatively, select or clear the Attach check box next to the name of each cluster to attach or detach the logical network to or from a given cluster.</p> <p>Required All - Allows you to specify whether the logical network is a required network on all clusters. Alternatively, select or clear the Required check box next to the name of each cluster to specify whether the logical network is a required network for a given cluster.</p>

5.3.5. Designate a Specific Traffic Type for a Logical Network with the Manage Networks Window

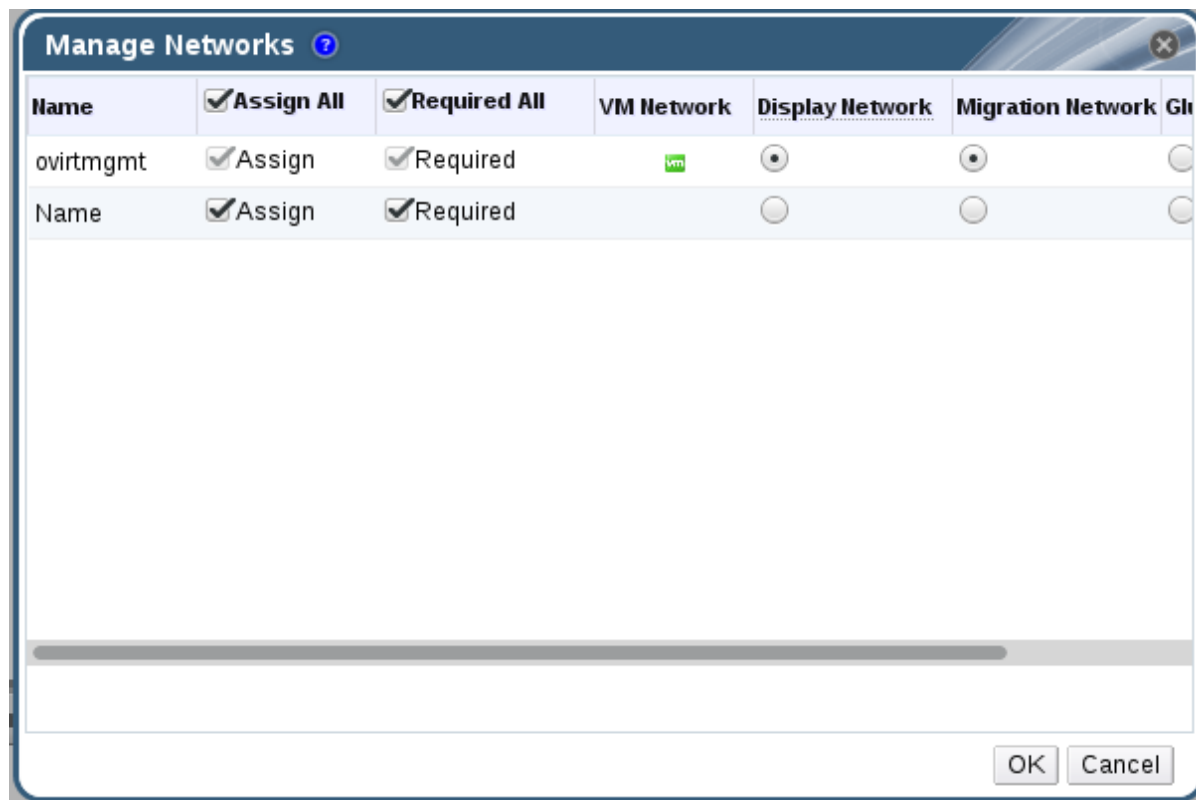
Summary

Specify the traffic type for the logical network to optimize the network traffic flow.

Procedure 5.3. Specifying Traffic Types for Logical Networks

Procedure 5.3. Specifying Traffic Types for Logical Networks

1. Click **Clusters** tab in tree mode and select the cluster in the results list.
2. Select the **Logical Networks** tab in the details pane to list the logical networks assigned to the cluster.
3. Click **Manage Networks** to open the **Manage Networks** window.

**Figure 5.5. Manage Networks Window**

4. Select appropriate check boxes.
5. Click **OK** to save the changes and close the window.

Result

You have optimized the network traffic flow by assigning a specific type of traffic to be carried on a specific logical network.

5.3.6. Explanation of Settings in the Manage Networks Window

The table below describes the settings for the **Manage Networks** window.

Table 5.3. Manage Networks Settings

Field	Description/Action
Assign	Assigns the logical network to all hosts in the cluster.

Field	Description/Action
Required	A Network marked "required" must remain operational in order for the hosts associated with it to function properly. If a required network ceases to function, any hosts associated with it become non-operational.
Gluster Network	A logical network marked "Gluster Network" carries gluster network traffic.

5.3.7. Network Labels

Network labels can be used to greatly simplify several administrative tasks associated with creating and administering logical networks and associating those logical networks with physical host network interfaces and bonds.

A network label is a plain text, human readable label that can be attached to a logical network or a physical host network interface. There is no strict limit on the length of label, but you must use a combination of lowercase and uppercase letters, underscores and hyphens; no spaces or special characters are allowed.

Attaching a label to a logical network or physical host network interface creates an association with other logical networks or physical host network interfaces to which the same label has been attached, as follows:

Network Label Associations

- When you attach a label to a logical network, that logical network will be automatically associated with any physical host network interfaces with the given label.
- When you attach a label to a physical host network interface, any logical networks with the given label will be automatically associated with that physical host network interface.
- Changing the label attached to a logical network or physical host network interface acts in the same way as removing a label and adding a new label. The association between related logical networks or physical host network interfaces is updated.

Network Labels and Clusters

- When a labeled logical network is added to a cluster and there is a physical host network interface in that cluster with the same label, the logical network is automatically added to that physical host network interface.
- When a labeled logical network is detached from a cluster and there is a physical host network interface in that cluster with the same label, the logical network is automatically detached from that physical host network interface.

Network Labels and Logical Networks With Roles

- When a labeled logical network is assigned to act as a display network or migration network, that logical network is then configured on the physical host network interface using DHCP so that the logical network can be assigned an IP address.

5.4. LOGICAL NETWORKS AND PERMISSIONS

5.4.1. Managing System Permissions for a Network

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them to a specific resource. For example **ClusterAdmin** has administrator privileges only for the assigned cluster.

A network administrator is a system administration role that can be applied for a specific network, or for all networks on a storage device, cluster or host. A network user can perform limited administration roles, such as viewing and attaching networks on a specific storage device or template. You can click the **Configure** button in the header bar to assign a network administrator for all networks in the environment.

The network administrator role permits the following actions:

- Create, edit and remove networks.
- Edit the configuration of the network, including configuring port mirroring.
- Attach and detach networks from resources including clusters and host.

The user who creates a network is automatically assigned **NetworkAdmin** permissions on the created network. You can also change the administrator of a network by removing the existing administrator and adding the new administrator.

5.4.2. Network Administrator and User Roles Explained

Network Permission Roles

The table below describes the administrator and user roles and privileges applicable to network administration.

Table 5.4. Red Hat Gluster Storage Network Administrator and User Roles

Role	Privileges	Notes
NetworkAdmin	Network Administrator for cluster or host. The user who creates a network is automatically assigned NetworkAdmin permissions on the created network.	Can configure and manage the network of a particular cluster or host. A network administrator of a cluster inherits network permissions for storage devices within the cluster. To configure port mirroring on a storage device network, apply the NetworkAdmin role on the network.
NetworkUser	Logical network and network interface user for virtual machine and template.	Can attach or detach network interfaces from specific logical networks.

CHAPTER 6. MANAGING RED HAT GLUSTER STORAGE HOSTS

A host is a physical, 64-bit server with an Intel or AMD chipset running Red Hat Gluster Storage 3.4. You can add new hosts to a storage cluster to expand the amount of available storage.

A host on the Red Hat Gluster Storage:

- Must belong to only one cluster in the system.
- Can have an assigned system administrator with system permissions.



IMPORTANT

Red Hat Gluster Storage Console uses various network ports for management. These ports must be open on Red Hat Gluster Storage hosts. For a full list of ports, see the *Red Hat Gluster Storage Console Installation Guide*.

6.1. HOSTS PROPERTIES

The **Hosts** tab provides a graphical view of all the hosts in the system.

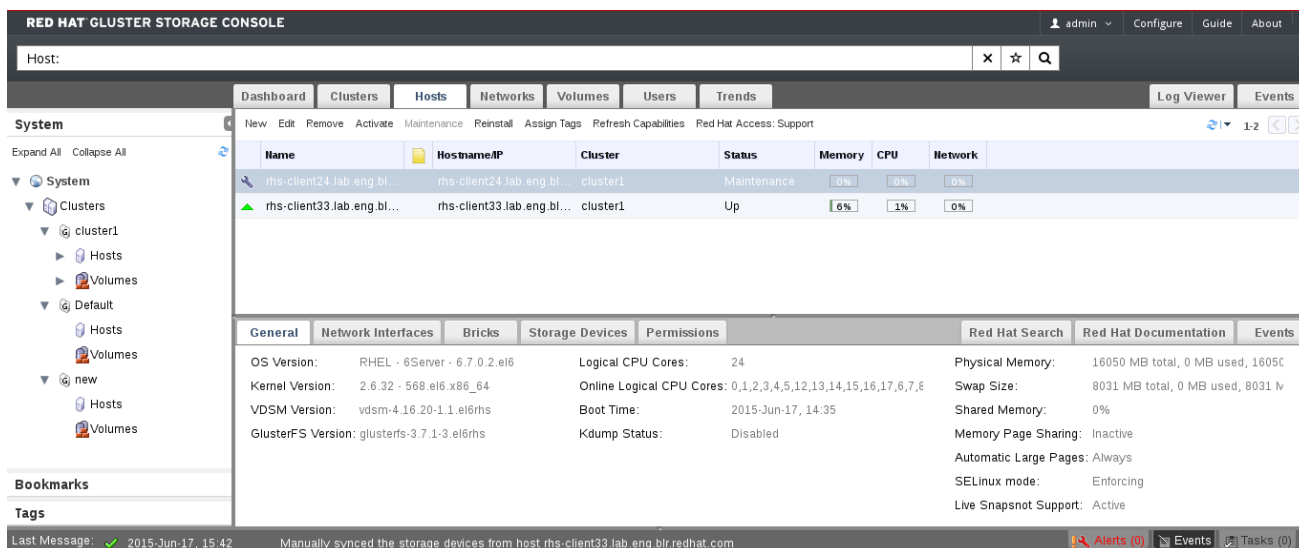


Figure 6.1. Hosts Details Pane

Table 6.1. Hosts Properties

Field	Description
Cluster	The selected cluster.
Name	The host name.
Address	The IP address or resolvable hostname of the host.

6.2. HOSTS OPERATIONS

6.2.1. Adding Hosts

You must install hosts and configure them with a name and IP address before you can add them to the Red Hat Gluster Storage Console.



IMPORTANT

If you re-install the Red Hat Gluster Storage Console, you must remove all hosts and reconnect them with the correct SSH keys for the new installation of Red Hat Gluster Storage Console.

Prerequisites

Before you can add a host to Red Hat Gluster Storage, ensure your environment meets the following criteria:

- The host hardware is Red Hat Enterprise Linux certified. See <https://access.redhat.com/ecosystem/#certifiedHardware> to confirm that the host has Red Hat certification.
- The host should have a resolvable hostname or static IP address.
- On Red Hat Enterprise Linux 7 nodes, register to the Red Hat Gluster Storage Server Channels if firewall needs to be configured automatically as iptables-service package is required.

Procedure 6.1. To Add a Host

Before adding a host, ensure you have the correct IP and password for the host. The process of adding a new host can take some time; you can follow its progress in the **Events** pane.

1. Click the **Hosts** tab to list available hosts.
2. Click **New** to open the **New Host** window.

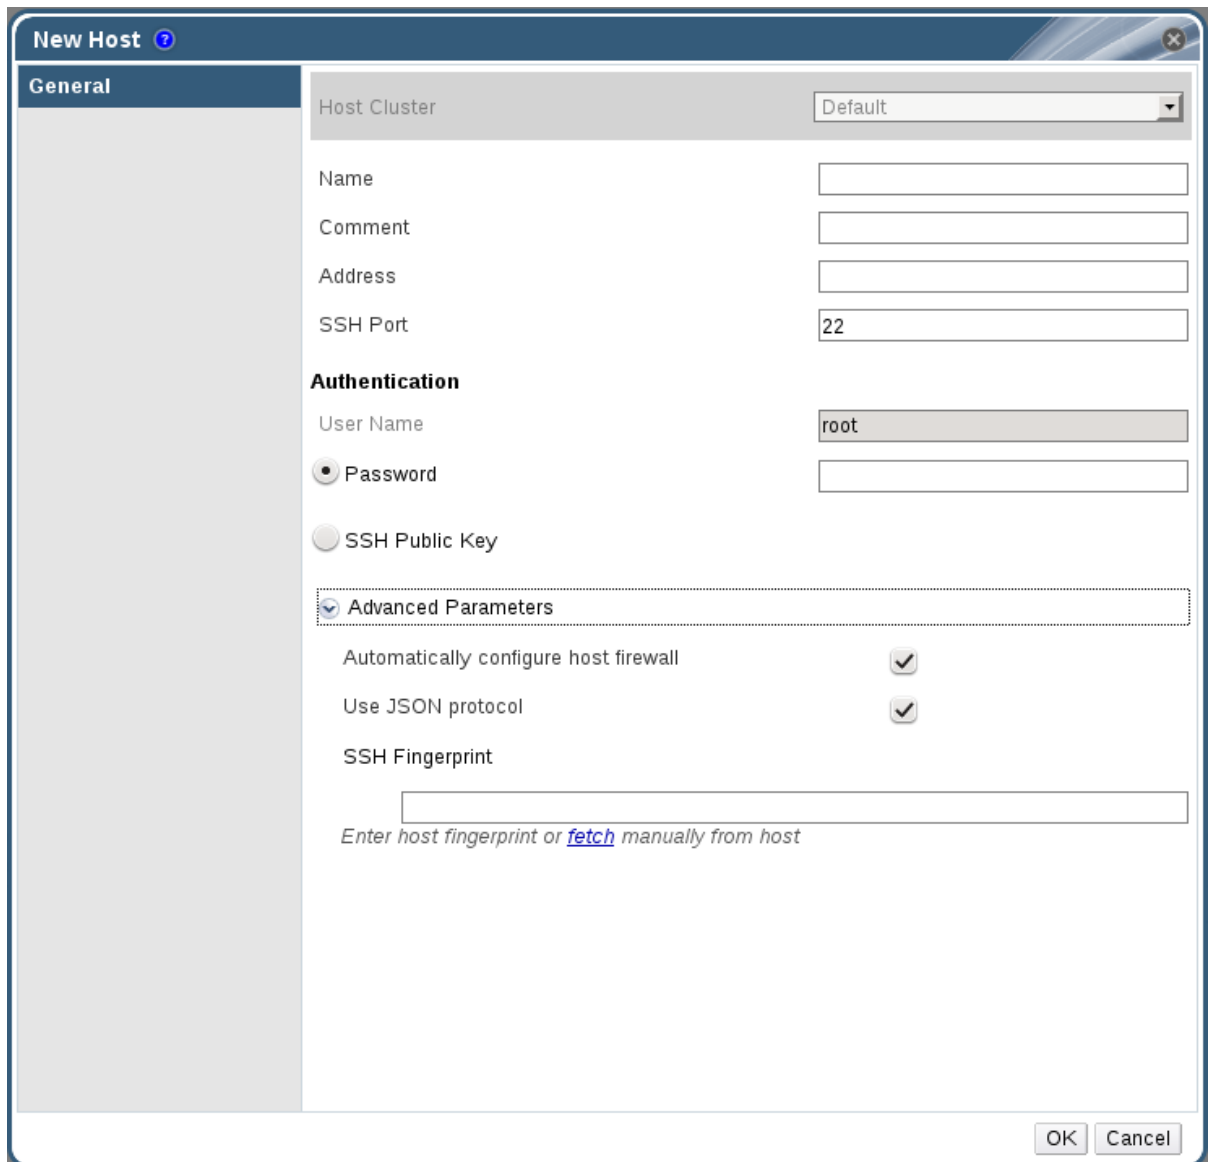


Figure 6.2. New Host Window

3. Select the **Host Cluster** for the new host from the drop-down menu.

Table 6.2. Add Hosts Properties

Field	Description
Host Cluster	The cluster to which the host belongs.
Name	<p>The name of the host. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.</p> <p>If Nagios is enabled, the host name given in Name field of Add Host window should match the host name given while configuring Nagios.</p>
Address	The IP address or resolvable hostname of the host.

Field	Description
Root Password	The password of the host's root user. This can only be given when you add the host, it cannot be edited afterwards.
SSH Public Key	Copy the contents in the text box to the <code>/root/.ssh/authorized_keys</code> file on the host if you'd like to use the Manager's ssh key instead of using a password to authenticate with the host.
Automatically configure host firewall	When adding a new host, the Manager can open the required ports on the host's firewall. This is enabled by default. This is an Advanced Parameter. The required ports are opened if this option is selected.
SSH Fingerprint	You can fetch the host's ssh fingerprint, and compare it with the fingerprint you expect the host to return, ensuring that they match. This is an Advanced Parameter.

**NOTE**

For Red Hat Enterprise Linux 7 hosts, iptables-service is used to manage firewall and existing firewalld configurations will not be enforced if "Automatically configure host firewall is chosen."

4. Enter the **Name**, and **Address** of the new host.
5. Select an authentication method to use with the host:
 - a. Enter the root user's password to use password authentication.
 - b. Copy the key displayed in the **SSH PublicKey** field to `/root/.ssh/authorized_keys` on the host to use public key authentication.
6. The mandatory steps for adding a Red Hat Gluster Storage host are complete. Click **Advanced Parameters** to show the advanced host settings:
 - a. Optionally disable automatic firewall configuration.
 - b. Optionally disable use of JSON protocol.

**NOTE**

With Red Hat Gluster Storage Console, the communication model between the engine and VDSM now uses JSON protocol, which reduces parsing time. As a result, the communication message format has changed from XML format to JSON format. Web requests have changed from synchronous HTTP requests to asynchronous TCP requests.

- c. Optionally add a host SSH fingerprint to increase security. You can add it manually, or fetch it automatically.

- Click **OK** to add the host and close the window.

The new host displays in the list of hosts with a status of "Installing" and then it goes to "Initialization" state and the host comes up.



NOTE

The host will be in *Up* status after the status of "Installing" and "Initialization" state. The host will have *Non-Operational* status when the host is not compatible with the cluster compatibility version. The *Non-Responsive* status will be displayed if the host is down or is unreachable.

You can view the progress of the host installation in the **Details** pane.

6.2.2. Activating Hosts

After taking down a host for maintenance, you must reactivate it before using it. When you activate a host, the host's networks are checked and glusterd service is restarted.

Procedure 6.2. To Activate a Host

- In the **Hosts** tab, select the host you want to activate.
- Click **Activate**. The host status changes to **Up**.

6.2.3. Managing Host Network Interfaces

The **Network Interfaces** tab in **Details** pane of **Hosts** enables you to attach a logical network to a host's physical network interface cards. The management network is defined, by default in the cluster. The administration portal automatically assigns the management network to the network interface which has the host address that was provided while adding a host.



NOTE

Interface with gluster network role should have an IP address configured. Ensure this by clicking Refresh Host Capabilities once host networks are setup.

6.2.3.1. Editing Host Network Interfaces

The **Network Interfaces** tab displays the name, network name, address, MAC address, speed, and link status for each interface. It also provides several options for managing host network interface cards.

Procedure 6.3. To Edit a Host Network Interface

- Click the **Hosts** tab to display a list of hosts. Select the desired host to display the **Details** pane.
- Click **Setup Host Networks** to open the **Setup Host Networks** window.




Figure 6.3. Setup Host Networks Window

3. Attach a logical network to a physical host network interface by selecting and dragging the logical network into the **Assigned Logical Networks** area next to the physical host network interface.
4. If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box. A logical network cannot be edited or moved to another interface until it is synchronized.

Alternatively, right-click the logical network and select a network interface from the drop-down menu.

5. Configure the logical network:
 1. Hover your cursor over an assigned logical network and click the pencil icon to open the **Edit Network** window.
 2. Select a **Boot Protocol**:
 - **None**
 - **DHCP**
 - **Static** - provide the **IP** and **Subnet Mask**.
6. Click **OK**.
7. Select **Verify connectivity between Host and Engine** to run a network check.
8. Select **Save network configuration** if you want the network changes to be persistent when you reboot the environment.

9. Click **OK** to implement the changes and close the window.

6.2.3.2. Editing Management/Gluster Network Interfaces

The **Network Interfaces** tab displays the name, network name, address, MAC address, speed, and link status for each interface. In the course of editing the host network interface cards, you may need to check or edit the management network interface. Similar procedure is applicable to Gluster network interface as well.



NOTE

Interface with gluster network role cannot be changed if gluster bricks are using it.



IMPORTANT

Clusters and hosts communicate through the management interface. Changing the properties of the management interface can cause the host to become unreachable.

Procedure 6.4. To Edit a Management Network Interface

1. Click the **Hosts** tab to display a list of hosts. Select the desired host to display the **Details** pane.
2. Edit the logical networks by hovering over an assigned logical network and clicking the pencil icon to open the **Edit Management Network** window.

Edit Management Network

Network Name: ovirtmgmt

Boot Protocol:

None

DHCP

Static

IP: 10.70.42.157

Subnet Mask: 255.255.252.0

Gateway: 10.70.43.254

Sync network

OK Cancel

Figure 6.4. Edit Management Network Dialog Box

3. If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box.
4. Select a **Boot Protocol**:
 - o **None**
 - o **DHCP**
 - o **Static** - provide the **IP** and **Subnet Mask**.
5. Make the required changes to the management network interface:
 - a. To attach the **ovirtmgmt** management network to a different network interface card, select a different interface from the **Interface** drop-down list.
 - b. Select the network setting from **None**, **DHCP** or **Static**. For the **Static** setting, provide the IP, Subnet and Default Gateway information for the host.
 - c. Click **OK** to confirm the changes.
 - d. Select **Verify connectivity between host and engine** if required.
 - e. Select **Save network configuration** to make the changes persistent when you reboot the environment.
6. Click **OK**.
7. Activate the host. See [Section 6.2.2, “Activating Hosts”](#).

6.2.4. Managing Gluster Sync

Gluster Sync periodically fetches the latest cluster configuration from glusterFS and synchronizes it with the engine database. The Red Hat Gluster Storage Console continuously monitors storage clusters for the addition and removal of hosts. If a change is detected, an action item displays in the **Cluster** tab with the option to **Import** or **Detach** the host.

Procedure 6.5. To Import a Host to a Cluster

1. Click the **Cluster** tab and select a cluster to display the **General** tab with details of the cluster.
2. In **Action Items**, click **Import** to display the **Add Hosts** window.

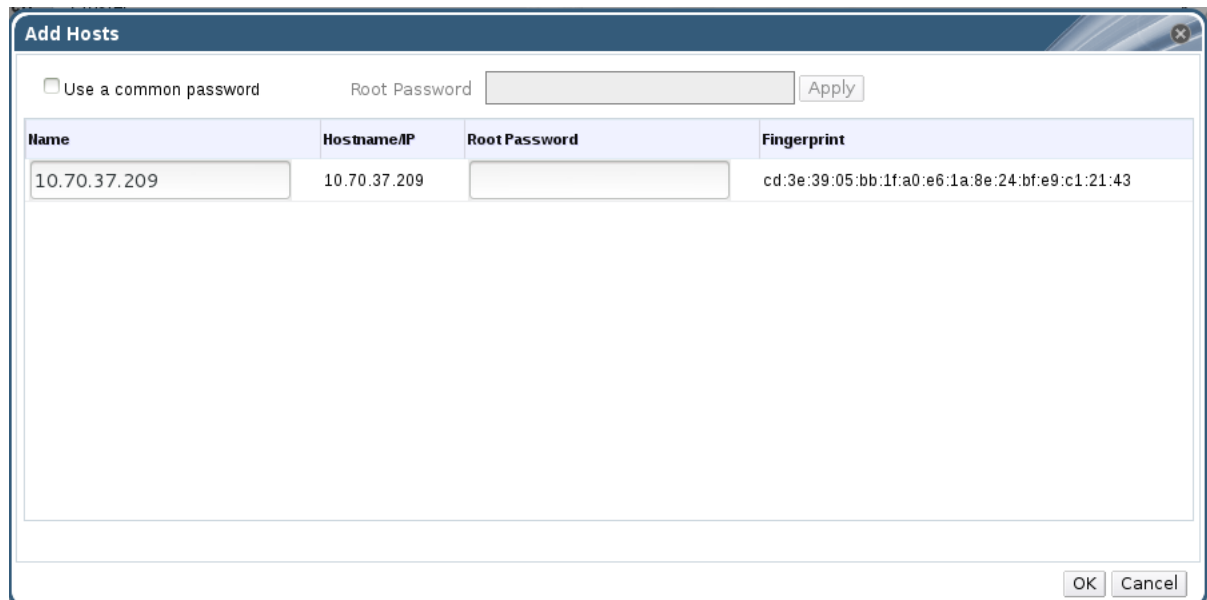


Figure 6.5. Add Hosts Window

3. Enter the **Name** and **Root Password**. Select **Use a common password** if you want to use the same password for all hosts.
4. Click **Apply**.
5. Click **OK** to add the host to the cluster.

Procedure 6.6. To Detach a Host from a Cluster

1. Click the **Cluster** tab and select a cluster to display the **General** tab with details of the cluster.
2. In **Action Items**, click **Detach** to display the **Detach Hosts** window.
3. Select the host you want to detach and click **OK**. Select **Force Detach** if you want to perform force removal of the host from the cluster.

6.2.5. Deleting Hosts

You can permanently remove hosts that are not in use. Deleting unused hosts saves system resources, as existing hosts are contacted at regular intervals.



NOTE

You can not remove hosts if it has volumes in it. Removing a host will detach the host from the cluster.

Procedure 6.7. To Delete a Host

1. Click the **Hosts** tab to display a list of hosts. Select the host you want to remove. If the required host is not visible, perform a search.
2. Click **Maintenance** to place the host into maintenance. Click **OK** to confirm the action. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance mode.

**IMPORTANT**

If you move the **Hosts** under *Maintenance* mode, it stops all gluster process such as brick, self-heal, and geo-replication. If you wish to reuse this host, ensure to remove the gluster related information stored in `/var/lib/glusterd` manually.

3. Click **Remove**.
4. Click **OK** to confirm.

6.2.6. Managing Storage Devices

You can view the list of storage devices and create bricks through Red Hat Gluster Storage Console.

**NOTE**

Click **Sync** to synchronize the storage devices from the selected host. By default, the synchronization will happen after every 2 hour.

**IMPORTANT**

RAID volumes must be provisioned externally before creating bricks. RAID configuration(RAID Type, No if disks, strip size) entered in the create brick dialog should match exactly with values used while creating RAID volume for better performance.

Procedure 6.8. Creating Bricks

1. Click the **Hosts** tab to display a list of hosts.
2. Select a host and select the **Storage Devices** sub-tab. The list of storage devices is displayed.
3. Select a storage device from the list and click **Create Brick**. The **Create Brick** page is displayed..
4. Enter the **Brick Name**, **Mount Point** name, and the **No. of Physical Disks in RAID Volume**.

The **Mount Point** is auto-suggested and can be edited.

5. Confirm the **Raid Type**.
6. Click **OK**. A new thinly provisioned logical volume is created with recommended Red Hat Gluster Storage configurations using the selected storage devices. This Logical Volume will be mounted at the specified mount point and this mount point can be used as brick in gluster volume.



IMPORTANT

Once the bricks are created from the UI and before using them to create a volume perform the following on all the nodes in the cluster:

- `semanage fcontext -a -t glusterd_brick_t '/rhgs/brick1(/.*)?'`
- `restorecon -Rv /rhgs/brick1`

Replace `/rhgs/brick1` with your actual brick path.

6.3. MAINTAINING HOSTS

You can use the Administration Portal to perform host maintenance tasks, for example changing the network configuration details of a host.




WARNING

Editing a host may require shutting down and restarting the host. Plan ahead when performing maintenance actions.

6.3.1. Moving Hosts into Maintenance Mode

To perform certain actions, you need to move hosts into maintenance mode.



IMPORTANT

If you move the Host under Maintenance mode, it will stop all gluster process such as brick, self-heal, and geo-replication. Once the Host is activated, all gluster processes will be restarted automatically.

Procedure 6.9. To Move a Host into Maintenance Mode

1. Click the **Hosts** tab to display a list of hosts.
2. Click **Maintenance** to place the host into maintenance. Click **OK** to confirm the action. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance mode.
3. Perform required tasks. When you are ready to reactivate the host, click **Activate**.
4. After the host reactivates, the **Status** field of the host changes to *Up*. If the Red Hat Gluster Storage Console is unable to contact or control the host, the **Status** field displays *Non-responsive*.

6.3.2. Editing Host Details


You can edit the details of a host, such as its name, network configuration, and cluster.

Procedure 6.10. To Edit Host Details

1. Click the **Hosts** tab to display a list of hosts.
2. If you are moving the host to a different cluster, first place it in maintenance mode by clicking **Maintenance**. Click **OK** to confirm the action. The **Status** field of the host changes to *Preparing for Maintenance*, followed by *Maintenance*. The icon changes to indicate that the host is in maintenance mode.
3. Click **Edit** to open the **Edit Host** dialog box.
4. To move the host to a different cluster, select the cluster from the **Host Cluster** drop-down list.
5. Make the required edits and click **OK**. Activate the host to start using it. See [Section 6.2.2, “Activating Hosts”](#)

6.3.3. Customizing Hosts

You can assign tags to help you organize hosts. For example, you can create a group of hosts running in a department or location.

**NOTE**

You can assign tags to a host only if a tag is present. You can not assign the *root* tag to a host. To create a new tag, see [Section 2.3, “Tags”](#).

Procedure 6.11. To Tag a Host

1. Click the **Hosts** tab to display a list of hosts. Select the desired host to display the **Details** pane.
2. Click **Assign Tags** to open the **Assign Tags** dialog box.
3. Select the required tags and click **OK**.

6.4. HOSTS ENTITIES**6.4.1. Viewing General Host Information**

The **General** tab on the **Details** pane provides information on individual hosts, including hardware and software versions, and available updates.

Procedure 6.12. To View General Host Information

1. Click the **Hosts** tab to display a list of hosts. If the required host is not visible, perform a search.
2. Select the desired host to display general information, network interface information and host information in the **Details** pane.
3. Click **General** to display the following information:
 - Version information for OS, Kernel, VDSM, and RHS.

- Status of memory page sharing (Active/Inactive) and automatic large pages (Always).
- CPU information: number of CPUs attached, CPU name and type, total physical memory allocated to the selected host, swap size, and shared memory.
- An alert if the host is in Non-Operational or Install-Failed state.

6.4.2. Viewing Network Interfaces on Hosts

The **Network Interfaces** tab on the **Details** pane provides information about the logical and physical networks on a host. This view enables you to define the attachment of the logical network in the Administration Portal to the physical network interface cards of the host. See [Section 6.2.3, “Managing Host Network Interfaces”](#) for more information on network interfaces.

Procedure 6.13. To View Network Interfaces on a Host

1. Click the **Hosts** tab to display a list of hosts. If the required host is not visible, perform a search.
2. Select the desired host to display the **Details** pane.
3. Click the **Network Interfaces** tab.

6.4.3. Viewing Permissions on Hosts

The **Permissions** tab on the **Details** pane provides information about user roles and their inherited permissions.

Procedure 6.14. To View Permissions on a Host

1. Click the **Hosts** tab to display a list of hosts. If the required host is not visible, perform a search.
2. Select the desired host to display the **Details** pane.
3. Click the **Permissions** tab.

6.4.4. Viewing Events from a Host

The **Events** tab on the **Details** pane provides information about important events such as notifications and errors.

Procedure 6.15. To View Events from a Host

1. Click the **Hosts** tab to display a list of hosts. If the required host is not visible, perform a search.
2. Select the desired host to display the **Details** pane.
3. Click the **Events** tab.

6.4.5. Viewing Bricks


The **Bricks** tab on the **Details** pane provides information about the bricks as the volume name and the brick directory.

Procedure 6.16. To View Bricks on a Host

1. Click the **Hosts** tab to display a list of hosts. If the required host is not visible, perform a search.
2. Select the desired host to display the **Details** pane.
3. Click the **Bricks** tab.

6.5. HOSTS PERMISSIONS

A host administrator has system administrator permissions for a specific host only. This role is useful when there are multiple hosts, each of which require their own system administrators. A host administrator has permissions for the assigned host only, not for all hosts in the cluster.

**NOTE**

You can only assign roles and permissions to existing users.

Procedure 6.17. To Add a Host Administrator Role

1. Click the **Hosts** tab to display a list of hosts. If the required host is not visible, perform a search.
2. Select the desired host to display the **Details** pane.
3. Click the **Permissions** tab to display a list of users and their current roles.

Name	Hostname/IP	Cluster	Status	Memory	CPU	Network
rhs-arch-srv1.lab.eng...	rhs-arch-srv1.lab.eng...	cluster2	Up	4%	0%	0%
rhs-client24.lab.eng.bl...	rhs-client24.lab.eng.bl...	cluster1	Up	7%	0%	0%
rhs-client33.lab.eng.bl...	rhs-client33.lab.eng.bl...	cluster1	Up	7%	0%	0%

User	Authorization provider	Namespace	Role
admin (admin@internal)	internal	*	SuperUser

Figure 6.6. Host Permissions Window

4. Click **Add** to display the **Add Permission to User** dialog box. Enter all or part of a name or user name in the **Search** box, then click **Go**. A list of possible matches displays in the results list.
5. Select the user you want to modify. Scroll through the **Role to Assign** list and select **HostAdmin**.
6. Click **OK** to display the name of the user and their assigned role in the **Permissions** tab.

Procedure 6.18. To Remove a Host Administrator Role

1. Click the **Hosts** tab to display a list of hosts. If the required host is not visible, perform a search.

2. Select the desired host to display the **Details** pane.
3. Click the **Permissions** tab to display a list of users and their current roles.
4. Select the desired user and click **Remove**


CHAPTER 7. MANAGING VOLUMES

You can use the console to create and start new volumes featuring a single global namespace. A volume is a logical collection of bricks where each brick is an export directory on a host in the trusted storage pool. Most of the management operations of Red Hat Gluster Storage Console happen on the volume.

A volume is the designated unit of administration in Red Hat Gluster Storage, so managing them is a large part of the administrator's duties.


The console also enables you to monitor the volumes in your cluster from the **Volumes** tab. To display the volumes, click the **Volumes** node from the **Tree** pane of the console window. The list of volumes is displayed in the right pane of the console window. It also displays the tasks and events for all volumes.

This chapter describes how to manage volumes stored on storage host machines.



NOTE

RDMA volumes can be managed through Red Hat Gluster Storage Console using the `glusternw` feature.



NOTE

Red Hat Gluster Storage Console does not support Dispersed volumes. If a Dispersed volume is created from the Gluster CLI, the details would be synchronized properly and listed in the Red Hat Gluster Storage Console. But users cannot manage Dispersed volumes from the Console.

7.1. CREATING A VOLUME

You can create new volumes in your storage environment. When creating a new volume, you must specify the bricks that comprise the volume and specify whether the volume is to be **Distribute**, **Replicate** or **Distributed Replicate**.

Procedure 7.1. Creating a Volume

1. Click the **Volumes** tab. The **Volumes** tab lists all volumes in the system.
2. Click the **New**. The **New Volume** window is displayed.

New Volume ?

Volume Cluster: Default

Name:

Type: Distribute

Transport Type: TCP RDMA

Bricks:
(0 bricks selected)

Access Protocols

Gluster:

NFS:

CIFS:


Allow Access From:
(Comma separated list of IP addresses/hostnames)

Optimize for Virt Store:

No host found in 'UP' state in the cluster, please select another cluster.

Figure 7.1. New Volume


3. Select the cluster from the **Volume Cluster** drop-down list.
4. In the **Name** field, enter the name of the volume.



NOTE


You can not create a volume with the name *volume*.

5. Select the type of the volume from the **Type** drop-down list. You can set the volume type to **Distribute**, **Replicate** or **Distributed Replicate**.

**NOTE**

- Creating replicated volumes with replica count more than 3 is under technology preview.

6. As necessary, click **Add Bricks** to add bricks to your volume.

**NOTE**

At least one brick is required to create a volume. The number of bricks required depends on the type of the volume.

For more information on adding bricks to a volume, see [Section 7.6.1, “Adding Bricks”](#).

7. Configure the **Access Protocol** for the new volume by selecting **NFS**, or **CIFS**, or both check boxes.
8. In the **Allow Access From** field, specify the volume access control as a comma-separated list of IP addresses or hostnames.

You can use wildcards to specify ranges of addresses. For example, an asterisk (*) specifies all IP addresses or hostnames. You need to use IP-based authentication for Gluster Filesystem and NFS exports.

You can optimize volumes for **virt-store** by selecting **Optimize for Virt Store**.

9. Click **OK** to create the volume. The new volume is added and displays on the **Volume** tab. The volume is configured, and **group** and **storage-owner-gid** options are set.

7.2. STARTING VOLUMES

After a volume has been created or an existing volume has been stopped, it needs to be started before it can be used.

Procedure 7.2. Starting a Volume

1. In the **Volumes** tab, select the volume to be started.

You can select multiple volumes to start by using the **Shift** or **Ctrl** key.

2. Click the **Start** button.

7.3. CONFIGURING VOLUME OPTIONS

Perform the following steps to configure volume options.

Procedure 7.3. Configuring Volume Options

1. Click the **Volumes** tab.

A list of volumes displays.

2. Select the volume to tune, and click the **Volume Options** tab from the **Details** pane.

The **Volume Options** tab lists the options set for the volume.

3. Click **Add** to set an option. The **Add Option** window is displayed. Select the option key from the drop-down list and enter the option value.

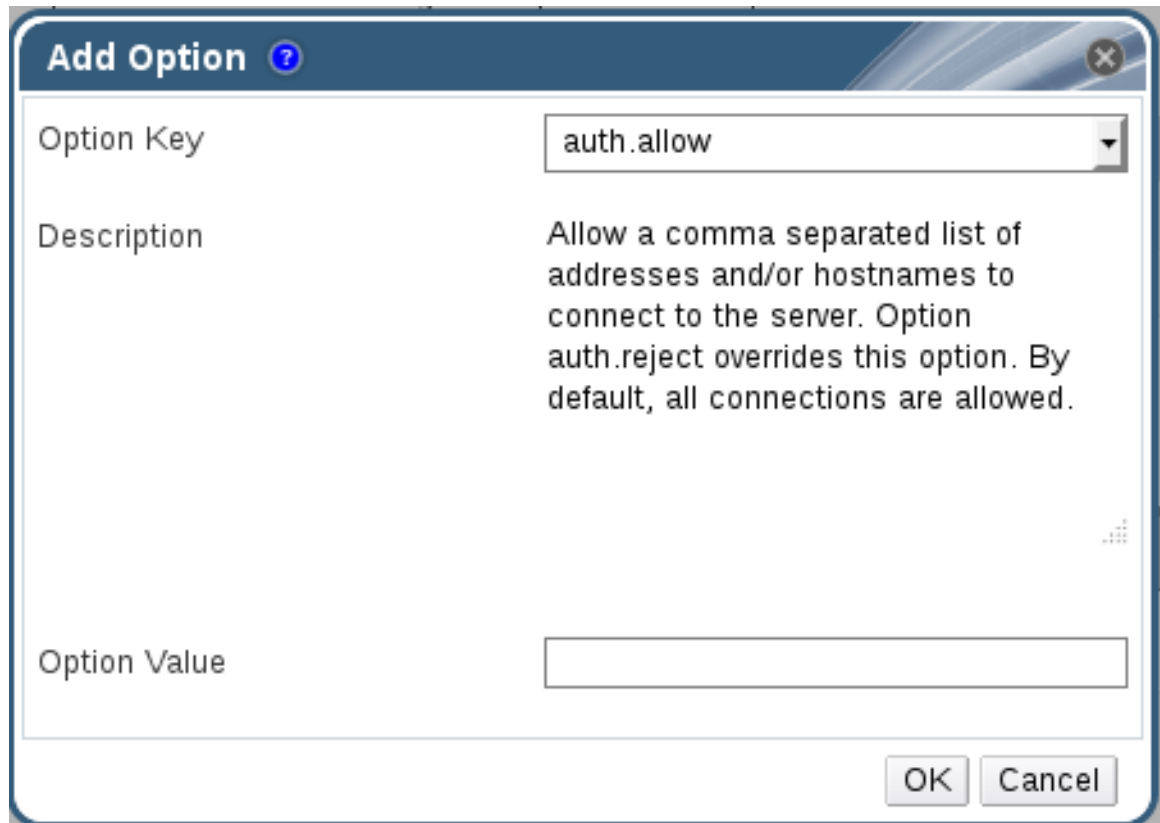


Figure 7.2. Add Option

4. Click **OK**.

The option is set and displays in the **Volume Options** tab.

For more information about volume options, see *Tuning Volume Options* in *Red Hat Gluster Storage Administration Guide*.

7.3.1. Edit Volume Options

Procedure 7.4. Editing Volume Options

1. Click the **Volumes** tab.

A list of volumes displays.

2. Select the volume to edit, and click the **Volume Options** tab from the **Details** pane.

The **Volume Options** tab lists the options set for the volume.

3. Select the option to edit. Click **Edit**. The **Edit Option** window is displayed. Enter a new value for the option in the **Option Value** field.

4. Click **OK**.

The edited option displays in the **Volume Options** tab.

7.3.2. Resetting Volume Options

Procedure 7.5. Resetting Volume Options

1. Click the **Volumes** tab.

A list of volumes is displayed.

2. Select the volume and click the **Volume Options** tab from the **Details** pane.

The **Volume Options** tab lists the options set for the volume.

3. Select the option to reset. Click **Reset**. **Reset Option** window is displayed, prompting to confirm the reset.

4. Click **OK**.

The selected option is reset. The name of the volume option reset is displayed in the **Events** tab.



NOTE

You can reset all volume options by clicking the **Reset All** button. A window is displayed, prompting to confirm the reset option. Click **OK**. All volume options are reset for the selected volume. A message that all volume options have been reset is displayed in the **Events** tab.

7.4. STOPPING VOLUMES

After a volume has been started, it can be stopped.



NOTE

You cannot stop a volume if there are any asynchronous tasks such as Rebalance or Remove Brick which are in progress.

Procedure 7.6. Stopping a Volume

1. In the **Volumes** tab, select the volume to be stopped.

You can select multiple volumes to stop by using the **Shift** or **Ctrl** key.

2. Click **Stop**. A window is displayed, prompting to confirm the stop.



NOTE

Stopping volume will make its data inaccessible.

3. Click **OK**.

7.5. DELETING VOLUMES

You can delete a volume or multiple volumes from your cluster.

Procedure 7.7. Deleting a Volume

1. In the **VoLumes** tab, select the volume to be deleted.
2. Click **Stop**. The volume stops.
3. Click **Remove**. A window is displayed, prompting to confirm the deletion. Click **OK**. The volume is removed from the cluster.

7.6. MANAGING BRICKS

A brick is the basic unit of storage, represented by an export directory on a host in the storage cluster. You can expand or shrink your cluster by adding new bricks or deleting existing bricks.



NOTE

When expanding distributed replicated volumes, the number of bricks being added must be a multiple of the replica count. For example, to expand a distributed replicated volume with a replica count of 2, you need to add bricks in multiples of 2 (such as 2, 4, 6, 8, etc.).

7.6.1. Adding Bricks

You can expand a volume by adding new bricks to an existing volume.

Procedure 7.8. Adding a Brick

1. Click the **VoLumes** tab.

A list of volumes displays.
2. Select the volume to which the new bricks are to be added. Click the **Bricks** tab from the **Details** pane.

The **Bricks** tab lists the bricks of the selected volume.
3. Click **Add** to add new bricks. The **Add Bricks** window is displayed.

Add Bricks ?

Volume Type Replica Count

Bricks

Host

Show available bricks from host

Brick Directory

Host	Brick Directory


Allow bricks in root partition and re-use the bricks by clearing xattrs

Figure 7.3. Add Bricks

Table 7.1. Add Bricks Tab Properties


Field/Tab	Description/Action
Volume Type	The type of volume.
Replica Count	Number of replicas to keep for each stored item.
Host	The selected host from which new bricks are to be added.
Brick Directory	The directory in the host.

4. Use the **Host** drop-down menu to select the host on which the brick resides.
5. Select the brick directory from the **Brick Directory** drop-down menu.


**NOTE**

Uncheck *Show available bricks from host* to type the brick directory path since the brick is not shown in the brick directory drop-down.

6. Select the *Allow bricks in root partition and re-use the bricks by clearing xattrs* to use the system's root partition for storage and to re-use the existing bricks by clearing the extended attributes.

**NOTE**

It is not recommended to reuse bricks of restored volume as is. In case of reusing the brick, delete the logical volume and recreate it from the same or different pool (the data on the logical volume will be lost). Otherwise there would be some performance penalty on copy-on-write because the original brick and the restored brick shares the block.

**NOTE**

Using the system's root partition for storage backend is not recommended. Original bricks of snapshot restored volume is not recommended to be used as a new brick

7. Click **Add** and click **OK**. The new bricks are added to the volume and is displayed in the **Bricks** tab.

7.6.2. Removing Bricks

You can shrink volumes as needed while the cluster is online and available. For example, to remove a brick that has become inaccessible in a distributed volume due to hardware or network failure.

**NOTE**

- When shrinking distributed replicated volumes, the number of bricks being removed must be a multiple of the replica count. For example, to shrink a distributed replicated volume with a replica count of 2, you need to remove bricks in multiples of 2 (such as 2, 4, 6, 8). In addition, the bricks you are removing must be from the same replica set. In a non-replicated volume, all bricks must be available in order to migrate data and perform the remove brick operation. In a replicated volume, at least one of the bricks in the replica must be available.
- You can monitor the status of *Remove Bricks* operation from the **Tasks** pane.
- You can perform *Commit*, *Retain*, *view Status* and *Stop* from remove-brick icon in the *Activities* column of *Volumes* and *Bricks* sub-tab.

Procedure 7.9. Removing Bricks from an Existing Volume

1. Click the **Volumes** tab.

A list of volumes is displayed.

2. Select the volume from which bricks are to be removed. Click the **Bricks** tab from the **Details** pane.

The **Bricks** tab lists the bricks for the volume.

3. Select the brick to remove. Click **Remove**. The **Remove Bricks** window is displayed, prompting to confirm the removal of the bricks.



WARNING

If the brick is removed without selecting the **Migrate Data from the bricks** check box, the data on the brick which is being removed will not be accessible on the glusterFS mount point. If the **Migrate Data from the bricks** check box is selected, the data is migrated to other bricks and on a successful commit, the information of the removed bricks is deleted from the volume configuration. Data can still be accessed directly from the brick.

4. Click **OK**, remove brick starts.



NOTE

- Once remove-brick starts, remove-brick icon is displayed in **Activities** column of both **Volumes** and **Bricks** sub-tab.
- After completion of the remove brick operation, the remove brick icon disappears after 10 minutes.

5. In the **Activities** column, ensure that data migration is completed and select the drop down of the remove-brick icon corresponding to the volume from which bricks are to be removed.

6. Click **Commit** to perform the remove brick operation.



Figure 7.4. Remove Bricks Commit


**NOTE**

The **Commit** option is enabled only if the data migration is completed.

The remove brick operation is completed and the status is displayed in the **Activities** column. You can check the status of the remove brick operation by selecting **Status** from the activities column.

7.6.2.1. Stopping a Remove Brick Operation

You can stop a remove brick operation after starting the remove brick operation. The remove brick operation is stopped and the migration of data is stopped.

**NOTE**

- Stop remove-brick operation is a technology preview feature. Technology Preview features are not fully supported under Red Hat subscription level agreements (SLAs), may not be functionally complete, and are not intended for production use. However, these features provide early access to upcoming product innovations, enabling customers to test functionality and provide feedback during the development process. As Red Hat considers making future iterations of Technology Preview features generally available, we will provide commercially reasonable efforts to resolve any reported issues that customers experience when using these features.
- Files which were migrated during Remove Brick operation are not migrated to the same brick when the operation is stopped.

Procedure 7.10. Stopping a Remove Brick Operation

1. Click the **Volumes** tab. A list of volumes displays.
2. In the **Activities** column, select the drop down of the remove-brick icon corresponding to the volume to stop remove brick.
3. Click **Stop** to stop remove brick operation. The remove brick operation is stopped and remove-brick icon in the activities column is updated. The remove brick status is displayed after stopping the remove brick.

You can also view the status of the *Remove Brick* operation by selecting **Status** from the drop down of the remove-brick icon in the *Activities* column of *Volumes* and *Bricks sub-tab*.

7.6.2.2. Viewing Remove Brick Status

You can view the status of a remove brick operation when the remove brick operation is in progress.

Procedure 7.11. Viewing Remove Brick Status

1. Click the **Volumes** tab. A list of volumes displays.
2. In the **Activities** column, click the arrow corresponding to the volume.
3. Click **Status** to view the status of the remove brick operation. The **Remove Bricks Status** window displays.

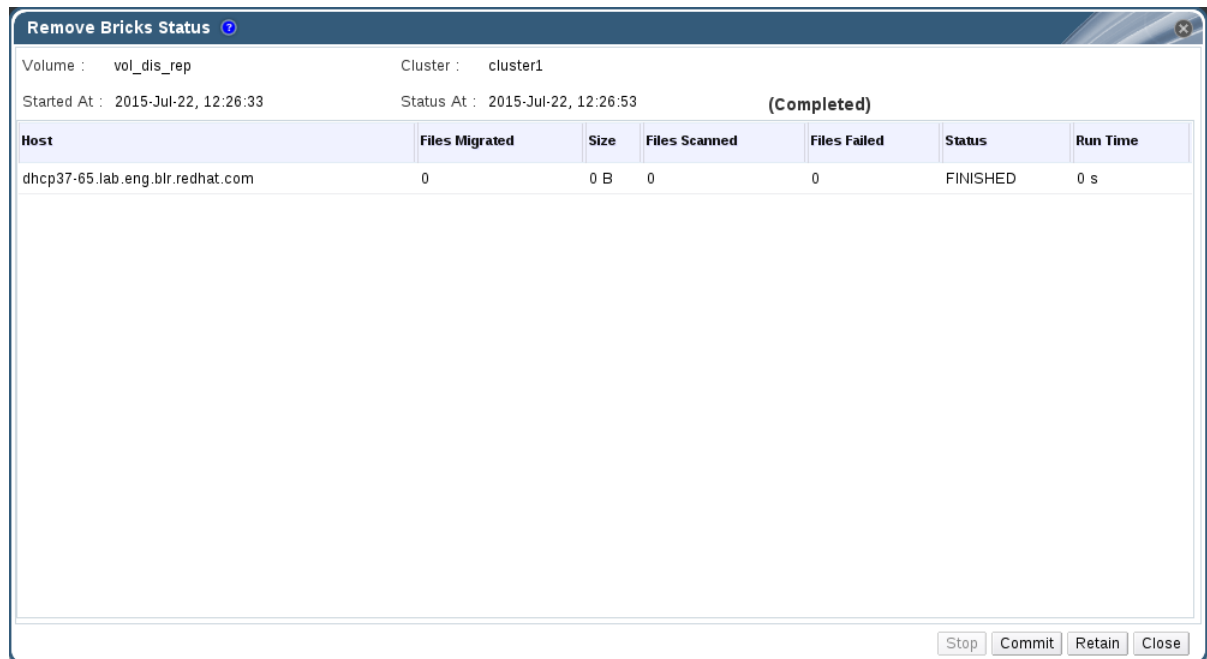


Figure 7.5. Remove Brick Status

4. Click one of the options below for the corresponding results
 - o **Stop** to stop the remove brick operation
 - o **Commit** to commit the remove brick operation
 - o **Retain** to retain the brick selected for removal
 - o **Close** to close the remove-brick status popup

7.6.2.3. Retaining a brick selected for Removal

You can retain a brick selected for removal operation when the remove brick operation is in progress. The brick that was selected to be removed will be retained and will not be removed from the volume.

The *Retain* option is enabled only after migration of data to other brick is completed.



NOTE

When a brick is retained, already migrated data is not migrated back.

Procedure 7.12. Retaining a Brick selected for Removal

1. Click the **Volumes** tab. A list of volumes displays.
2. In the **Activities** column, click the arrow corresponding to the volume.
3. Click **Retain** to retain the brick selected for removal. The brick is not removed and the status of the operation is displayed in the remove brick icon in the *Activities* column.

You can also check the status by selecting the **Status** option from the drop down of remove-brick icon in the activities column.

7.6.3. Viewing Advanced Details

You can view the advanced details of a particular brick of the volume through Red Hat Gluster Storage Console. The advanced view displays the details of the brick, and is divided into four parts. Namely, **General**, **Clients**, **Memory Statistics**, and **Memory Pools**.

Procedure 7.13. Viewing Advanced Details

1. Click the **Volumes** tab. A list of volumes displays.
2. Select the required volume and click the **Bricks** tab from the **Details** pane.
3. Select the brick and click **Advanced Details**. The **Brick Advanced Details** window displays.

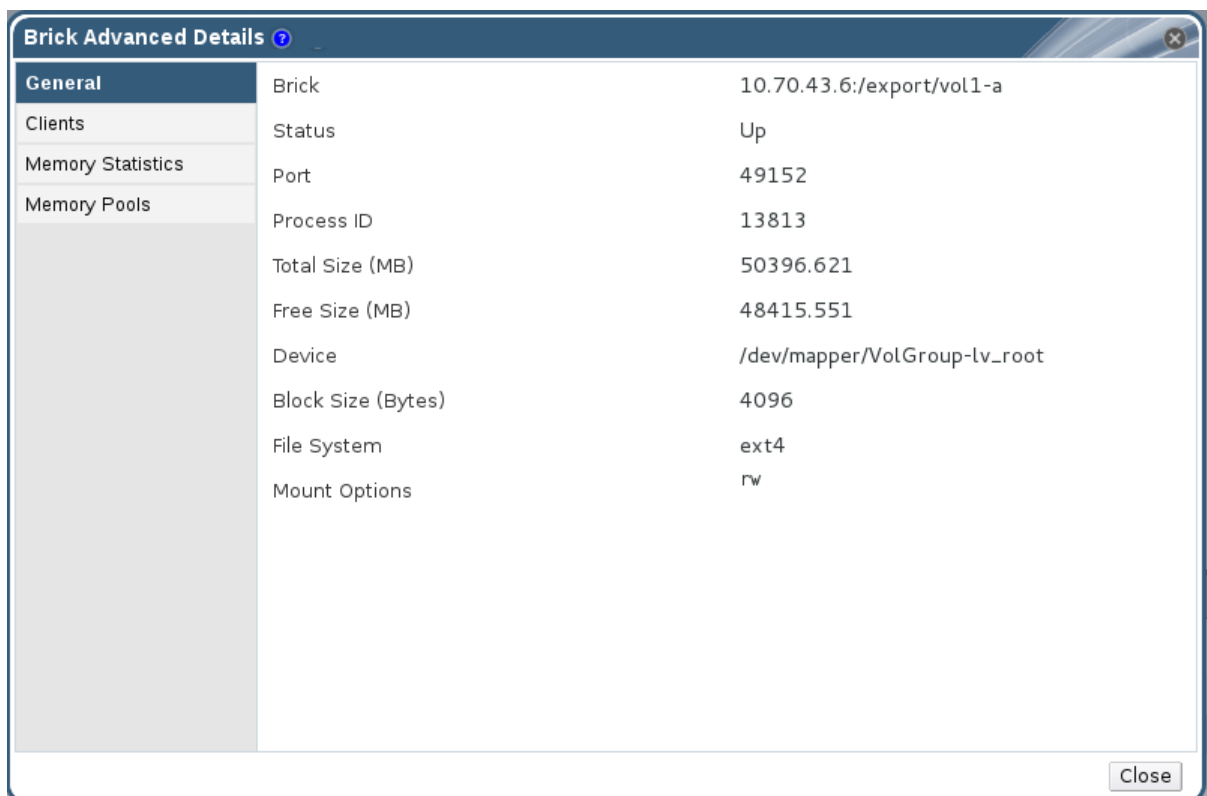


Figure 7.6. Brick Advanced Details

Table 7.2. Brick Details

Field/Tab	Description/Action
General	Displays additional information about the bricks.
Clients	Displays a list of clients accessing the volumes.
Memory Statistics/Memory Pool	Displays the details of memory usage and memory pool for the bricks.

You can view the advanced details of Red Hat Gluster Storage volumes through Red Hat Gluster Storage Console.

7.7. VOLUMES PERMISSIONS

While the superuser or system administrator of the Red Hat Gluster Storage has the full range of permissions, a Storage Administrator is a system administration role for a specific volume only. This is a hierarchical model, meaning that the Cluster Administrator has permissions to manage volumes. However, Storage Administrators have permission for the assigned volumes only, and not for all volumes in the cluster.

Procedure 7.14. Assigning a System Administrator Role for a Volume

1. Click the **Volumes** tab. A list of volumes displays.
2. Select the volume to edit, and click the **Permissions** tab from the **Details** pane.

The **Permissions** tab lists users and their current roles and permissions, if any.

The screenshot shows the Red Hat Gluster Storage Console interface. The top navigation bar includes 'Dashboard', 'Clusters', 'Hosts', 'Networks', 'Volumes', 'Users', and 'Trends'. The 'Volumes' tab is active, displaying a table of volumes:

Name	Cluster	Volume Type	Bricks	Info	Space Used	Activities	No of snapshots
gluster_shared_storage	cluster1	Distributed Replicates	4	0	2%		0
vol1	cluster1	Distribute	2	0	0%		0
vol2	cluster1	Distribute	6	0	0%		1

Below the table, the 'Permissions' sub-tab is selected, showing a table of users:

User	Authorization provider	Namespace	Role
admin (admin@internal)	internal	*	SuperUser

The interface also includes a sidebar with 'System', 'Clusters', 'Hosts', and 'Volumes' sections, and a bottom status bar showing the last message: '2 brick(s) added to volume gluster_shared_storage of cluster cluster1.'

Figure 7.7. Volume Permissions

3. Click **Add** to add an existing user. The **Add Permission to User** window is displayed. Enter a name, a user name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches displays in the results list.
4. Select the check box of the user to be assigned the permissions. Scroll through the **Role to Assign** list and select **GlusterAdmin**.

Figure 7.8. Assign GlusterAdmin Permission

5. Click **OK**.

The name of the user displays in the **Permissions** tab, with an icon and the assigned role.



NOTE

You can only assign roles and permissions to existing users.

You can also change the system administrator of a volume by removing the existing system administrator and adding a new system administrator as described in the previous procedure.

Procedure 7.15. Removing a System Administrator Role

1. Click the **Volumes** tab. A list of volumes displays.
2. Select the required volume and click the **Permissions** tab from the **Details** pane.

The **Permissions** tab lists users and their current roles and permissions, if any. The Super User and Cluster Administrator, if any, will display in the **Inherited Permissions** tab. However, none of these higher level roles can be removed.

3. Select the appropriate user.
4. Click **Remove**. A window is displayed, prompting to confirm removing the user. Click **OK**. The user is removed from the **Permissions** tab.

7.8. REBALANCING VOLUME

Storage volumes are abstracted from hardware, allowing each to be managed independently. Storage can be added or removed from the storage pools while data continues to be available, with no application interruption. Volumes can be expanded or shrunk across machines by adding or removing bricks.

After expanding or shrinking a volume (without migrating data), you need to rebalance the data among the hosts. In a non-replicated volume, all bricks should be online to perform the rebalance operation. In a replicated volume, at least one of the bricks in the replica should be online.

Through Red Hat Gluster Storage Console, you can perform the following:

- Start Rebalance
- Stop Rebalance
- View Rebalance Status



NOTE

You can monitor the status of Rebalance operation from the **Tasks** pane.

7.8.1. Start Rebalance

1. Click the **Volumes** tab. The **Volumes** tab is displayed with the list of all volumes in the system.
2. Select the volume that you want to **Rebalance**.
3. Click the **Rebalance**. The **Rebalance** process starts and the rebalance icon is displayed in the **Activities** column of the volume. A mouseover script is displayed mentioning that the rebalance is in progress. You can view the rebalance status by selecting status from the rebalance drop-down list .



NOTE

After completion of the rebalance operation, the rebalance icon disappears after 10 minutes.

7.8.2. Stop Rebalance

1. Click the **Volumes** tab. The **Volumes** tab is displayed with the list of all volumes in the system.
2. Select a volume on which Rebalance need to be stopped.



NOTE


- You can not stop rebalance for multiple volumes.
- Rebalance can be stopped for volumes only if it is in progress

3. In the **Activities** column, select the drop-down of the Rebalance icon corresponding to the volume.
4. Click **Stop**. The **Stop Rebalance** window is displayed.
5. Click **OK** to stop rebalance. The **Rebalance** is stopped and the status window is displayed.

You can also check the status of the *Rebalance* operation by selecting **Status** option from the drop down of Rebalance icon in the activities column.

7.8.3. View Rebalance Status

1. Click the **Volumes** tab. The **Volumes** tab is displayed with the list of all volumes in the system.
2. Select the volume on which Rebalance is in progress, stopped, completed.
3. Click **Status** option from the Rebalance icon drop down list. The **Rebalance Status** page is displayed.



The screenshot shows a window titled "Rebalance Status" with the following information:

- Volume : gluster_shared_storage
- Cluster : cluster1
- Started At : 2015-Jun-19, 12:48:12
- Status At : 2015-Jun-19, 12:48:46
- (Completed)

Host	Files Rebalanced	Size	Files Scanned	Files Failed	Files Skipped	Status	Run Time
rhs-client24.lab.eng.blr.redhat.com	0	0 B	3	0	0	FINISHED	0 s
rhs-client33.lab.eng.blr.redhat.com	0	0 B	3	0	3	FINISHED	0 s

Buttons: Stop, Close

Figure 7.9. Rebalance Status



NOTE

If the Rebalance Status window is open while Rebalance is stopped using the CLI, the status is displayed as **Stopped**. If the **Rebalance Status** window is not open, the task status is displayed as **Unknown** as the status update depends on gluster CLI.

You can also stop *Rebalance* operation by clicking **Stop** in the **Rebalance Status** window.


CHAPTER 8. MANAGING GLUSTER HOOKS

Gluster hooks are volume lifecycle extensions. You can manage Gluster hooks from Red Hat Gluster Storage Console. The content of the hook can be viewed if the hook content type is **Text**. Through Red Hat Gluster Storage Console, you can perform the following:

- View a list of hooks available in the hosts.
- View the content and status of hooks.
- Enable or disable hooks.
- Resolve hook conflicts.

8.1. VIEWING THE LIST OF HOOKS

Click the **Cluster** tab and select a cluster. A **Gluster Hooks** sub-tab displays, listing the hooks in the cluster.



The screenshot shows the Red Hat Gluster Storage Console interface. The main content area displays a table of Gluster Hooks for a selected cluster. The table has the following data:


Name	Status	Volume Event	Stage	Content Type
28Quota-enable-root-xattr-heal...	Enabled	add-brick	Pre	Text
57glusterfind-delete-post.py	Enabled	delete	Post	Text
57glusterfind-delete-post.pyc	Enabled	delete	Post	Binary
57glusterfind-delete-post.pyo	Enabled	delete	Post	Binary
56glusterd-geo-rep-create-post...	Enabled	gsync-create	Post	Text

Figure 8.1. Gluster Hooks

8.2. VIEWING THE CONTENT OF HOOKS

Procedure 8.1. Viewing the Content of a Hook

1. Click the **Cluster** tab and select a cluster. A **Gluster Hooks** sub-tab displays, listing the hooks in the cluster.
2. Select a hook with content type **Text** and click **View Content**. The **Hook Content** window displays with the content of the hook.

A screenshot of a window titled "Hook Content" with a question mark icon. The window contains a text area with shell script comments. The text is as follows:

```
#!/bin/bash

#Need to be copied to hooks/<HOOKS_VER>/set/post/

#TODO: All gluster and samba paths are assumed for fedora like systems.
#Some efforts are required to make it work on other distros.

#The preferred way of creating a smb share of a gluster volume has changed.
#The old method was to create a fuse mount of the volume and share the mount
#point through samba.
#
#New method eliminates the requirement of fuse mount and changes in fstab.
#glusterfs_vfs plugin for samba makes call to libgfapi to access the volume.
#
#This hook script enables user to enable or disable smb share by volume set
```

A "Close" button is visible in the bottom right corner of the window.

Figure 8.2. Hook Content

8.3. ENABLING OR DISABLING HOOKS

Procedure 8.2. Enabling or Disabling a Hook

1. Click the **Cluster** tab and select a cluster. A **Gluster Hooks** sub-tab displays, listing the hooks in the cluster.
2. Select the hook and click **Enable** or **Disable**.

If **Disable** is selected, *Disable Gluster Hooks* dialog box displays, prompting you to confirm disabling hook. Click *OK* to confirm disabling.

The hook is enabled or disabled on all nodes of the cluster.

The enabled or disabled hooks status update displays in the **Gluster Hooks** sub-tab.

8.4. REFRESHING HOOKS

By default, the Red Hat Gluster Storage Console checks for the status of installed hooks on all hosts in the cluster and detects new hooks by running a periodic job every hour. If the user wishes to trigger this job, they can choose to do so by clicking **Sync**.

Procedure 8.3. Refreshing a Hook

1. Click the **Cluster** tab and select a cluster. A **Gluster Hooks** sub-tab displays, listing the hooks in the cluster.
2. Click **Sync**. The hooks are synchronized and displayed.

8.5. RESOLVING CONFLICTS

The hooks are displayed in the **Gluster Hooks** sub-tab of the **Cluster** tab. Hooks causing a conflict are displayed with an exclamation mark. This denotes either that there is a conflict in the content or status of the hook across the servers in the cluster, or that the hook script is missing in one or more servers. These conflicts can be resolved via the console. The hooks in the servers are periodically synchronized with engine database and the following conflicts can occur for the hooks:

- Content Conflict - the content of the hook is different across servers.
- Status Conflict - the status of the hook is different across servers.
- Missing Conflict - one or more servers of the cluster do not have the hook.
- Content + Status Conflict - both the content and status of the hook are different across servers.
- Content + Status + Missing Conflict - both the content and status of the hook are different across servers, or one or more servers of the cluster do not have the hook.

8.5.1. Resolving Missing Hook Conflicts

Procedure 8.4. Resolving a Missing Hook Conflict

1. Click the **Cluster** tab and select a cluster. A **Gluster Hooks** sub-tab displays, listing the hooks in the cluster.
2. Select a hook causing a conflict and click **Resolve Conflicts**. The **Resolve Conflicts** window displays.

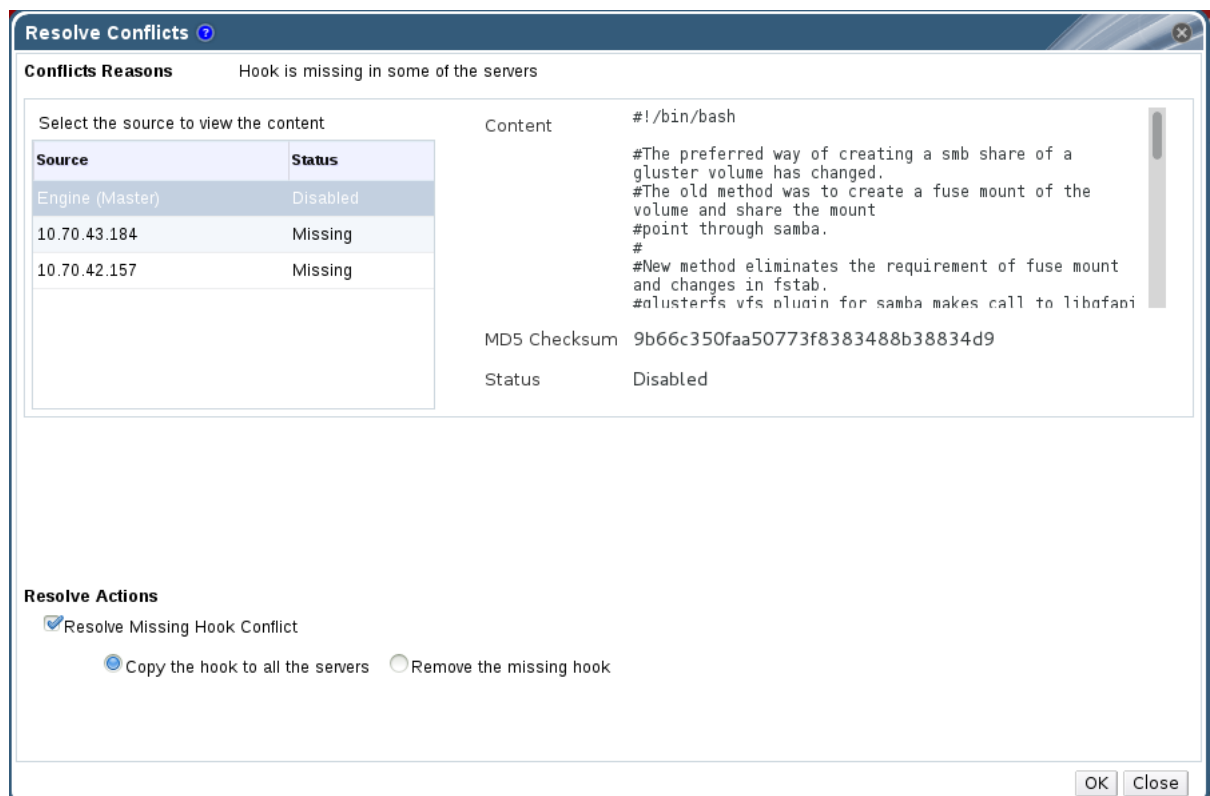


Figure 8.3. Missing Hook Conflict

3. Select one of the options give below:

- o **Copy the hook to all the servers** to copy the hook to all servers.
 - o **Remove the missing hook** to remove the hook from all servers and the engine.
4. Click **OK**. The conflict is resolved.

8.5.2. Resolving Content Conflicts

Procedure 8.5. Resolving a Content Conflict

1. Click the **Cluster** tab and select a cluster. A **Gluster Hooks** sub-tab displays, listing the hooks in the cluster.
2. Select the conflicted hook and click **Resolve Conflicts**. The **Resolve Conflicts** window displays.

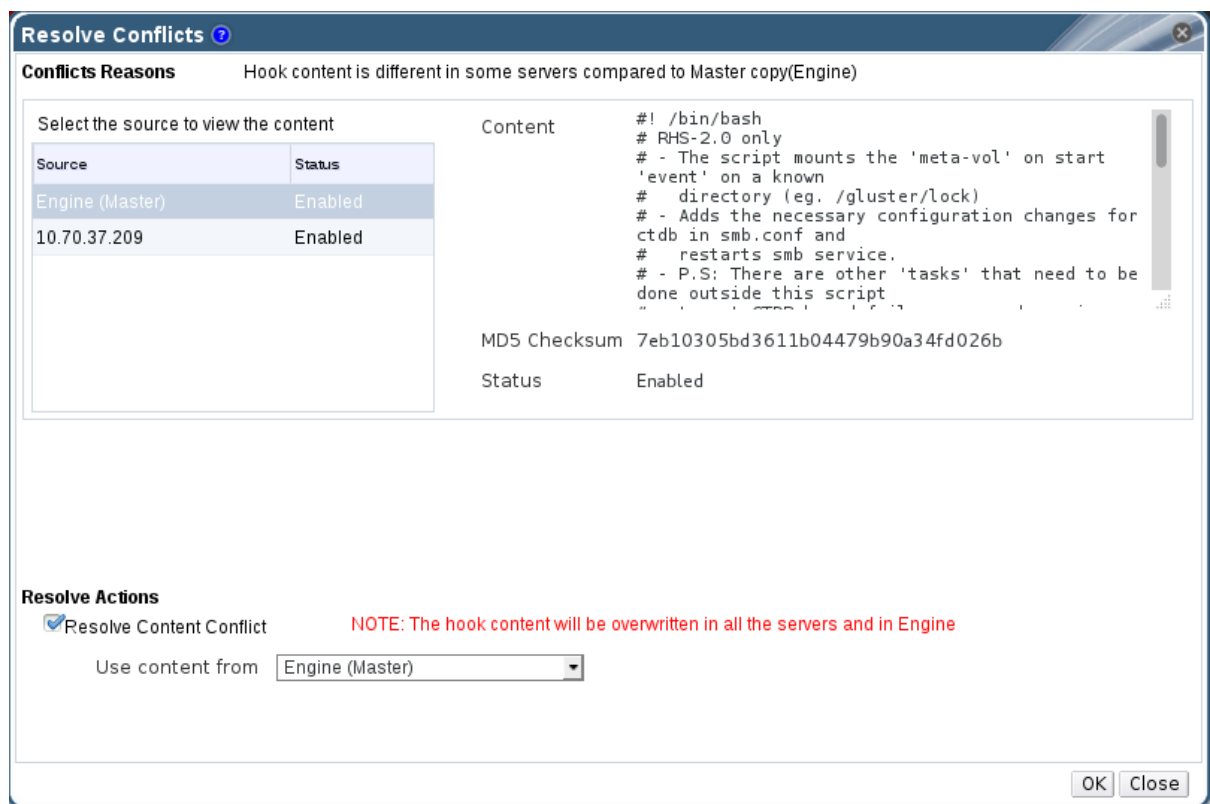


Figure 8.4. Content Conflict

3. Select an option from the **Use Content from** drop-down list:
 - o Select a server to copy the content of the hook from the selected server.
 - Or
 - o Select **Engine (Master)** to copy the content of the hook from the engine copy.



NOTE

The content of the hook will be overwritten in all servers and in the engine.

4. Click **OK**. The conflict is resolved.

8.5.3. Resolving Status Conflicts

Procedure 8.6. Resolving a Status Conflict

1. Click the **Cluster** tab and select a cluster. A **Gluster Hooks** sub-tab displays, listing the hooks in the cluster.
2. Select the conflicted hook and click **Resolve Conflicts**. The **Resolve Conflicts** window displays.

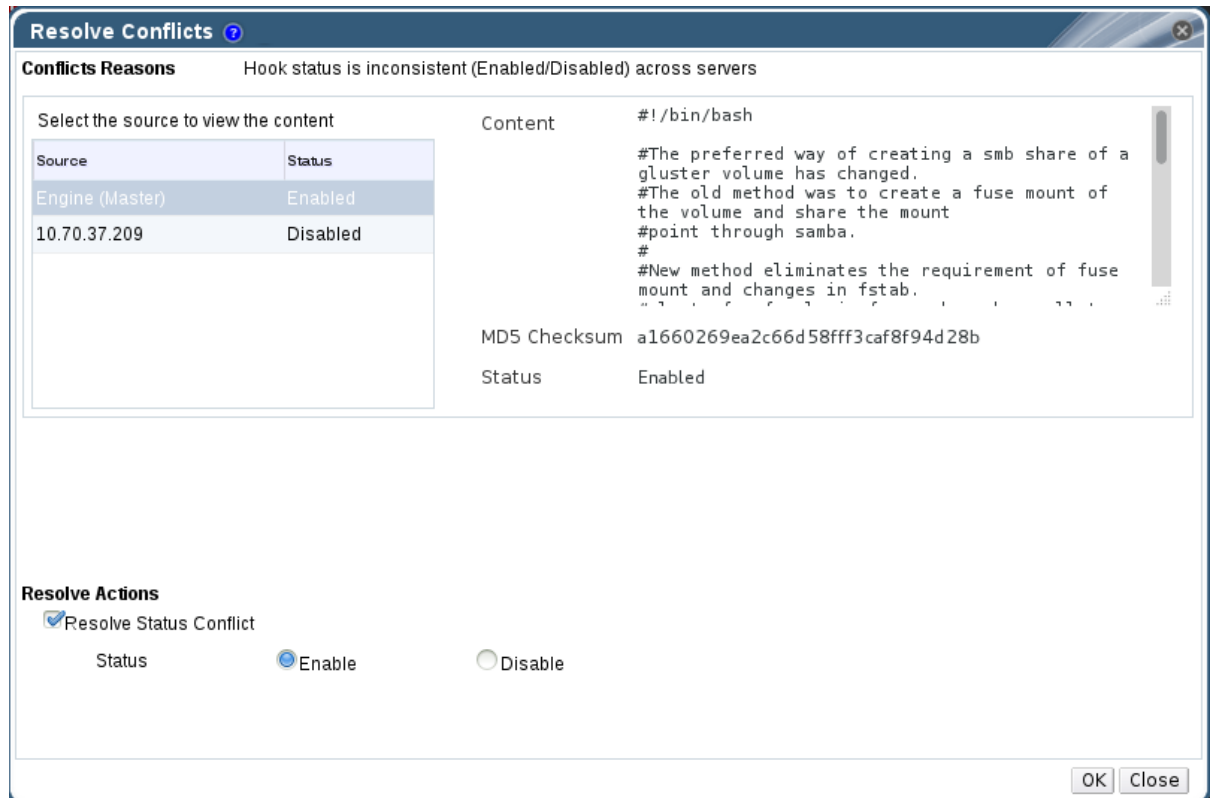


Figure 8.5. Status Conflict

3. Set **Hook Status** to **Enable** or **Disable**.
4. Click **OK**. The conflict is resolved.


8.5.4. Resolving Content and Status Conflicts

Procedure 8.7. Resolving a Content and Status Conflict

1. Click the **Cluster** tab and select a cluster. A **Gluster Hooks** sub-tab displays, listing the hooks in the cluster.
2. Select a hook causing a conflict and click **Resolve Conflicts**. The **Resolve Conflicts** window displays.
3. Select an option from the **Use Content from** drop-down list to resolve the content conflict:
 - o Select a server to copy the content of the hook from the selected server.

Or

- Select **Engine (Master)** to copy the content of the hook from the engine copy.

**NOTE**

The content of the hook will be overwritten in all the servers and in engine.

4. Set **Hook Status** to **Enable** or **Disable** to resolve the status conflict.
5. Click **OK**. The conflict is resolved.


8.5.5. Resolving Content, Status, and Missing Conflicts

Procedure 8.8. Resolving a Content, Status and Missing Conflict

1. Click the **Cluster** tab and select a cluster. A **Cluster Hooks** sub-tab displays, listing the hooks in the cluster.
2. Select the conflicted hook and click **Resolve Conflicts**. The **Resolve Conflicts** window displays.
3. Select one of the options given below to resolve the missing conflict:
 - Copy the hook to all the servers.
 - Remove the missing hook.
4. Select an option from the **Use Content from** drop-down list to resolve the content conflict:
 - Select a server to copy the content of the hook from the selected server.

Or

 - Select **Engine (Master)** to copy the content of the hook from the engine copy.

**NOTE**

The content of the hook will be overwritten in all the servers and in Engine.

5. Set **Hook Status** to **Enable** or **Disable** to resolve the status conflict.
6. Click **OK**. The conflict is resolved.

CHAPTER 9. MANAGING SNAPSHOTS

Red Hat Gluster Storage Console Snapshot feature enables you to create point-in-time copies of Red Hat Gluster Storage volumes, which you can use to protect data. You can directly access read-only Snapshot copies to recover from accidental deletion or modification of the data. Through Red Hat Gluster Storage Console, you can view the list of snapshots and snapshot status, create, delete, activate, deactivate and restore to a given snapshot.


For more information on the highlights of Snapshot feature and prerequisites, refer *Managing Snapshots* chapter of *Red Hat Gluster Storage Administration Guide*.

9.1. CREATING SNAPSHOTS

You can create snapshots of Red Hat Gluster Storage Volumes to preserve its contents at a single point in time. This enables you to recover data from accidental deletion or modification.

Procedure 9.1. Creating Snapshots

1. Click the **VoLumes** tab. The list of all volumes is displayed.
2. Select the volume of which you want to create Snapshot.
3. Click **Snapshot** and click **New** to open the **Create Snapshot** page.



General	Cluster	cluster_RHEL6.7
Schedule	Volume	vol1
	Snapshot Name Prefix	<input type="text"/>
	Description	<input type="text"/>

Figure 9.1. Creating Snapshots

4. Enter the **Snapshot Name Prefix** and **Description**.
5. Click **OK** to create Snapshot.

The format of the snapshot is <Snapshot name prefix>_<Timezone of RHS node>-<yyyy>.<MM>.<dd>-<hh>.<mm>.<ss>

For information on scheduling a snapshot creation, see [Section 9.3, “Scheduling Snapshots”](#).

9.2. CONFIGURING SNAPSHOTS

You can set the following configuration parameters related to Snapshot for a specific volume or cluster:

Clusters

- **Hard Limit:** If the snapshot count in a volume reaches this limit then no further snapshot creation is allowed. The range is from 1 to 256. Once this limit is reached you have to remove the snapshots to create further snapshots. This limit can be set for the system or per volume. If both system limit and volume limit is configured then the effective max limit would be the lowest of the two value.
- **Soft limit:** This is a percentage value. The default value is 90%. This configuration works along with auto-delete feature. If auto-delete is enabled then it will delete the oldest snapshot when snapshot count in a volume crosses this limit. When auto-delete is disabled it will not delete any snapshot, but it will display a warning message to the user.
- **Auto deletion flag:** This will enable or disable auto-delete feature. By default auto-delete is disabled. When enabled it will delete the oldest snapshot when snapshot count in a volume crosses the snap-max-soft-limit. When disabled it will not delete any snapshot, but it will display a warning message to the user.
- **Activate-on-Create:** Volume snapshots would be auto activated after creation.

Volumes

Hard Limit: If the snapshot count in a volume reaches this limit then no further snapshot creation is allowed. The range is from 1 to 256. Once this limit is reached you have to remove the snapshots to create further snapshots. This limit can be set for the system or per volume. If both system limit and volume limit is configured then the effective max limit would be the lowest of the two value.

Procedure 9.2. Configuring Snapshots

1. Click the **Volumes** tab. The list of all volumes in the system is displayed.
2. Select the volume for which you want to configure Snapshot.

If a volume is not selected from the list, only the cluster level parameters can be modified and set.

3. Click **Snapshot**.
4. Click **Options - Clusters** or **Options - Volume** to configure Snapshot for Cluster or Volume respectively.

Configuring Snapshot configurations for Cluster

If a volume is not selected, the cluster level parameters can be modified.

1. Click **Snapshot** and select **Options- Clusters**.

2. Select the cluster from the drop down list.
3. Modify the **Snapshot Options**. You can set the hard limit, soft limit percentage and enable or disable auto deletion of Snapshots for Clusters.
4. Click Update to update the details.

Configuring Snapshot configurations for Volume

1. Click **Snapshot** and select **Options- Volume**.
2. Modify the **Snapshot Options**. You can set maximum number of Snapshots for selected volume.
3. Click Update to update the details.

For information on Configuring Snapshot behavior see, *Configuring Snapshot behavior* section in the *Red Hat Gluster Storage Administration Guide*.

9.3. SCHEDULING SNAPSHOTS

When a snapshot is scheduled, an icon is displayed in the info column of the volumes main tab which shows that snapshot scheduling is done. Once the schedule expires, the icon will be removed from the info column. You can schedule the snapshot creation for the selected volume using (recurrence type) option. By default, the recurrence type value is set to *None*.



NOTE

You can delete existing Snapshot schedule by changing the recurrence type to *None*.

Procedure 9.3. Scheduling Snapshots

1. Click the **Volumes** tab. The list of all volumes in the system is displayed.
2. Select the volume of which you want to schedule Snapshot.
3. Click **Snapshot** and click **New** to open **Create/Schedule Snapshot** page.
4. In **General** tab, enter **Snapshot Name Prefix** and **Description**.
5. Click **Schedule** tab.

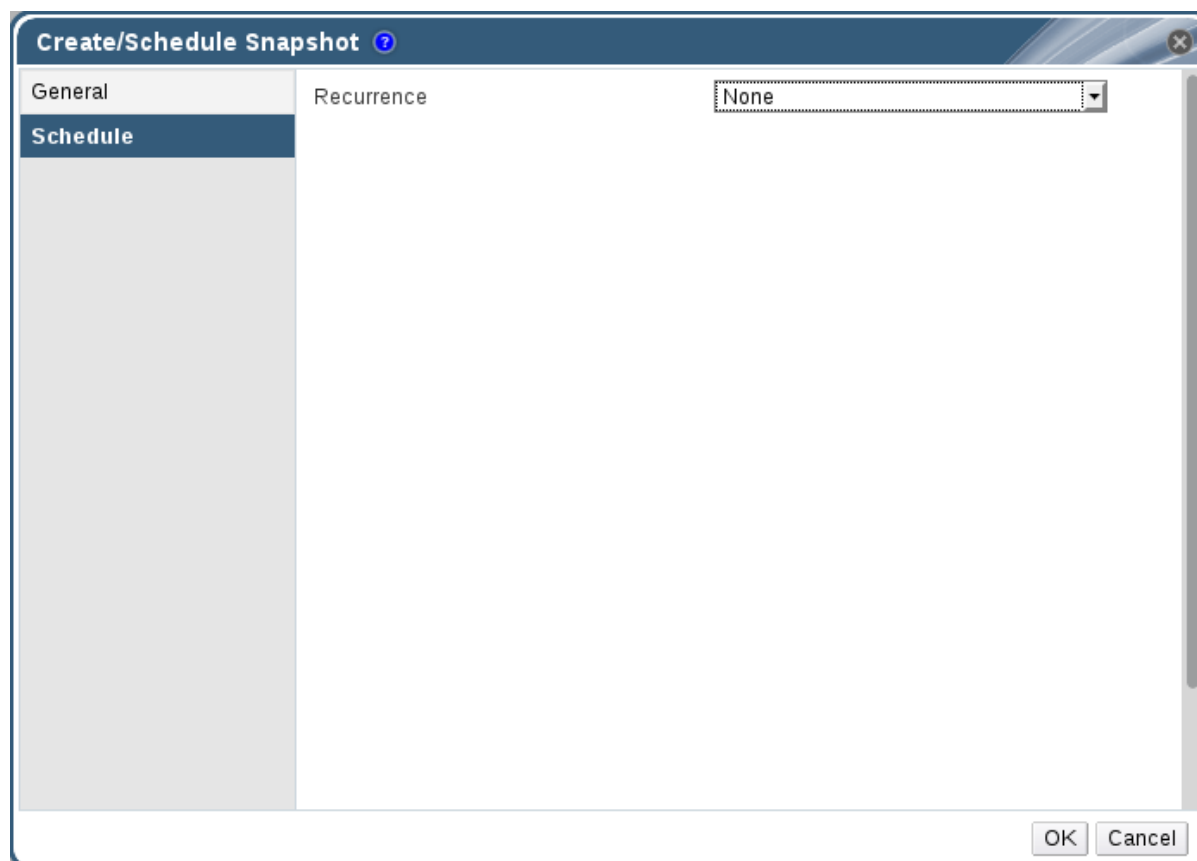


Figure 9.2. Scheduling Snapshots

6. Select the recurrence schedule for the Snapshot. You can schedule the snapshot to recur at intervals of a specified number of minutes, hours, days, weeks, or months; either perpetually, or between specified dates.

Create/Schedule Snapshot

General

Schedule

Recurrence: Minute

Interval (every n-mins): 5

Time Zone: Atlantic/Reykjavik=(GMT) Greenwich

Start Schedule by: 27-Jul-2015 15:50

End By: Date No End Date

End Schedule by: 27-Jul-2015 15:50

Frequent creation of snapshots would overload the cluster

Gluster CLI based snapshot scheduling is enabled. It would be disabled once volume snapshots scheduled from UI.

OK Cancel

Figure 9.3. Recurrence Schedule

Set **Recurrence** to the unit of time that you want to use as an interval between snapshots. If you do not want to set up recurring snapshots, leave this field set to **None**.

Minutes

Takes a snapshot every N minutes, where N is the value of the **Interval** field, with the first snapshot being taken at the time specified in the **Start Schedule by** field.

Hours

Takes a snapshot every N hours, where N is the value of the **Interval** field, with the first snapshot being taken after the time specified in the **Start Schedule by** field. Subsequent snapshots will be taken at the start of the hour. For example, if snapshots are to recur every 2 hours, and the first snapshot occurs at 2.20 PM, the next snapshot will occur at 4.00 PM.

Days

Takes a snapshot at the time specified in the **Execute At** field every N days, where N is the value of the **Interval** field, with the first snapshot being taken after the time specified in the **Start Schedule by** field.

Weeks

Takes a snapshot at the time specified in the **Execute At** field every N weeks, where N is the value of the **Interval** field, with the first snapshot being taken after the time specified in the **Start Schedule by** field.


Months

Takes a snapshot at the time specified in the **Execute At** field every N months, where N is the value of the **Interval** field, with the first snapshot being taken at the time specified in the **Start Schedule by** field.

The **End by** option determines whether snapshots will stop after a certain date. To set an end date, set **End by** to **Date**, and use the fields beside **End Schedule By** to enter a date and time at which snapshots should stop. To take snapshots continuously with no end date, set **End by** to **No End Date**.


7. Click OK to set the snapshot recurrence schedules.

If a soft limit or a hard limit for the number of snapshots of a volume reaches a limit, an alert is generated.



NOTE

Syncing of snapshots takes 5 minutes from the CLI.



IMPORTANT


Scheduling of snapshots cannot be done from CLI once it is scheduled from Red Hat Gluster Storage Console. In the case of Red Hat Gluster Storage Nodes are no longer managed by Red Hat Gluster Storage Console, snapshot scheduling could be done using gluster CLI as well by running the command `echo "none" > /var/run/gluster/shared_storage/snaps/current_scheduler`.

9.4. RESTORING SNAPSHOTS

You can restore Red Hat Gluster Storage volume to the state of a specific snapshot.

Procedure 9.4. Restoring Snapshots

1. Click the **Volumes** tab. The list of all volumes in the system is displayed.
2. Select the volume for which you want to restore the Snapshot.
3. Click **Snapshots** sub-tab and select the Snapshot.
4. Click **Restore** and click **OK** to confirm Snapshot restore.



NOTE

While restoring a snapshot, you will lose current state and the volume will be brought down and restored to the state of the selected snapshot.

9.5. ACTIVATING SNAPSHOTS

To use a snapshot, you must make the snapshots online by activating snapshot.

Procedure 9.5. Activating Snapshots

1. Click the **Volumes** tab. The list of all volumes in the system is displayed.

2. Select the volume for which you want to activate the Snapshot.
3. Click **Snapshots** sub-tab and select the Snapshot.
4. Click **Activate** and click **OK** to confirm activation of the Snapshot.

9.6. DEACTIVATING SNAPSHOTS

You can deactivate an active snapshot of a volume if it is not required to be in active state anymore.

Procedure 9.6. Deactivating Snapshots

1. Click the **Volumes** tab. The list of all volumes in the system is displayed.
2. Select the volume for which you want to deactivate the Snapshot.
3. Click **Snapshots** sub-tab and select the Snapshot.
4. Click **Deactivate** and click **OK** to confirm Snapshot deactivation.

9.7. DELETING SNAPSHOTS

You can delete existing snapshot if the snapshot will not be used any further.

Procedure 9.7. Deleting Snapshots

1. Click the **Volumes** tab. The list of all volumes in the system is displayed.
2. Select the volume for which you want to delete the Snapshot.
3. Click **Snapshots** sub-tab and select the Snapshot.
4. Click **Delete** and click **OK** confirm deleting the selected Snapshot.


To delete all Snapshots for the selected volume, click **Delete All**.

CHAPTER 10. MANAGING GEO-REPLICATION

Geo-replication provides a distributed, continuous, asynchronous, and incremental replication service from one site to another over Local Area Networks (LANs), Wide Area Networks (WANs), and the Internet.

Geo-replication uses a source-destination model, where replication and mirroring occurs between the following partners:

- Source – a Red Hat Gluster Storage volume.
- Destination - a Red Hat Gluster Storage volume.




NOTE

Red Hat Gluster Storage recommends the usage of meta-volume for geo-replication sessions. For information on meta-volume see *Configuring a Meta-Volume* section in *Red Hat Gluster Storage Administration Guide*.

10.1. GEO-REPLICATION OPERATIONS

You can perform Geo-replication operations and also manage source and destination volumes through Red Hat Gluster Storage Console. The Info column in Volumes tab displays icon which indicates the presence of a Geo-replication session and also if the volume is source or a destination volume. Hovering over the icon also provides information about if the volume is source or destination.



IMPORTANT

- Manually set the cluster option "cluster.enable-shared-storage" from CLI.
- Set the option use_meta_volume to true.
- For every new node added to the cluster ensure that the cluster option "cluster.enable-shared-storage" is set to the cluster and the meta-volume is mounted.

10.1.1. Creating a Geo-replication session

You can create geo-replication sessions for volumes eligible for geo-replication. You can select "Show volumes eligible for geo-replication" option, to view the list of volumes eligible for geo-replication. If the "Show volumes eligible for geo-replication" option is not selected, all the volumes are displayed. Geo-replication session creation may fail if you select volumes which are not eligible for geo-replication session.


Destination volume eligibility criteria

- Destination and source volumes should not be from same cluster.
- Capacity of destination volume should be greater than or equal to that of source volume.
- Cluster Compatibility version of destination and source volumes should be same.
- Destination volume should not already be a part of another geo replication session.
- Destination volume should be up.

- Destination volume should be empty.

Procedure 10.1. Creating a Geo-replication session

1. Click the **Volumes** tab. The list of volumes in the system is displayed.
2. Select the volume for which the geo-replication is to be created and click the **Geo-replication** option.
3. Click **New** option. The **New Geo-Replication Session** page is displayed.



NOTE

- You can also create Geo-Replication session from Geo-Replication sub-tab.

Figure 10.1. New Geo-replication session

4. Select the **Destination Cluster**, **Destination Volume**, and **Destination Host**.
5. Select **Show volumes eligible for geo-replication** option to view the list of volumes eligible for geo-replication.
6. Enter the **User Name**. For non-root user, enter the corresponding **User group**.

7. Select **Auto-start geo-replication session after creation** option to start the session immediately after creation and click **OK**.

10.1.2. Viewing Geo-replication session Details

You can view the list of all the individual nodes, their status, Up time, and also the detailed status for each node.

The status of a Geo-replication session can be one of the following:

- **Initializing:** This is the initial phase of the Geo-replication session; it remains in this state for a minute in order to make sure no abnormalities are present.
- **Created:** The geo-replication session is created, but not started.
- **Active:** The gsync daemon in this node is active and syncing the data.
- **Passive:** A replica pair of the active node. The data synchronization is handled by active node. Hence, this node does not sync any data.
- **Faulty:** The geo-replication session has experienced a problem, and the issue needs to be investigated further.
- **Stopped:** The geo-replication session has stopped, but has not been deleted.
- **Crawl Status**
 - **Changelog Crawl:** The changelog translator has produced the changelog and that is being consumed by gsyncd daemon to sync data.
 - **Hybrid Crawl:** The gsyncd daemon is crawling the glusterFS file system and generating pseudo changelog to sync data.
- **Checkpoint Status:** Displays the status of the checkpoint, if set. Otherwise, it displays as N/A.

Procedure 10.2. Viewing a Geo-replication session Details

To view the geo-replication session details, follow the steps given below:

1. Click the **Volumes** tab. The list of volumes is displayed.
2. Select the desired volume and click **Geo-Replication** sub-tab.
3. Select the session from the **Geo-Replication** sub-tab.
4. Click **View Details**. The Geo-replication details, Destination Host, Destination Volume, User Name, and Status are displayed.

10.1.3. Starting or Stopping a Geo-replication session



IMPORTANT

You must create the geo-replication session before starting geo-replication. For more information, see [Section 10.1.1, “Creating a Geo-replication session”](#)

**NOTE**

Stopping Geo-replication session will fail if:

- any node that is a part of the volume is offline.
- if it is unable to stop the geo-replication session on any particular node.
- if the geo-replication session between the master and slave is not active.

To start geo-replication, follow the steps given below:

Procedure 10.3. Starting and Stopping Geo-replication session

1. Click the **Volumes** tab. The list of volumes is displayed.
2. Select the desired volume and click the **Geo-Replication** sub-tab.
3. Select the session from the **Geo-Replication** sub-tab.
4. Click **Start** or **Stop** to start or stop the session respectively.

**NOTE**

Click **Force start session** to force the operation on geo-replication session on the nodes that are part of the master volume. If it is unable to successfully perform the operation on any node which is online and part of the master volume, the command will still perform the operation on as many nodes as it can. This command can also be used to re-perform the operation on the nodes where the session has died, or the operation has not been executed.

10.1.4. Pausing or Resuming a Geo-replication session**Procedure 10.4. Pausing or Resuming Geo-replication session**

1. Click the **Volumes** tab. The list of volumes is displayed.
2. Select the desired volume and click the **Geo-Replication** sub-tab.
3. Select the session from the **Geo-Replication** sub-tab.
4. Click **Pause** or **Resume** to pause or resume the Geo-replication session.

10.1.5. Removing a Geo-replication session**NOTE**

You must first stop a geo-replication session before it can be deleted. For more information, see *Starting or Stopping Geo-replication session* section.

Procedure 10.5. Removing Geo-replication session

To remove/delete a geo-replication session, follow the steps given below:


1. Click the **Volumes** tab. The list of volumes is displayed.
2. Select the desired volume and click the **Geo-Replication** sub-tab.
3. Select the session from the **Geo-Replication** sub-tab.
4. Click **Remove** to remove the Geo-replication session.

10.1.6. Synchronizing a Geo-replication session

Procedure 10.6. Synchronizing a Geo-replication session

To synchronize a geo-replication session, follow the steps given below:

1. Click the **Volumes** tab. The list of volumes is displayed.
2. Select the desired volume and click the **Geo-Replication** sub-tab.
3. Select the session from the **Geo-Replication** sub-tab.
4. Click **Sync**. The geo-replication session is synchronized.



NOTE

By default:

- New Sessions and Config options for a session are synced every 1 hour.
- Existing sessions status are updated every 5 minutes.
- Session configs are also automatically synchronized whenever a config is set/reset from the Console.

10.1.7. Configuring Options for a Geo-replication

You can change the values of the geo-replication options at any point of time after creating the geo-replication session. The geo-replication session will be restarted automatically if the user changes any configuration.

Procedure 10.7. Configuring Options for a Geo-replication

To configure options for geo-replication, follow the steps given below:

1. Click the **Volumes** tab. The list of volumes is displayed.
2. Select the desired volume and click the **Geo-Replication** sub-tab.
3. Select the session from the **Geo-Replication** sub-tab.
4. Click **Options**. The **Geo-Replication Options** page is displayed.
5. To set the config, modify the **Option Value**. To reset, select reset check box corresponding to the **Option Key**.
6. Click **OK**. The **Option Keys** are modified for the Geo-Replication session.

10.1.8. Non-root Geo-replication

Geo-replication supports access to Red Hat Gluster Storage slaves through SSH using an unprivileged account (user account with non-zero UID). This method is recommended as it is more secure and it reduces the master's capabilities over slave to the minimum. This feature relies on mountbroker, an internal service of glusterd which manages the mounts for unprivileged slave accounts.

Prerequisites

- Ensure to create the geo-replication users and groups using CLI. Currently, creation of users and groups using Red Hat Gluster Storage Console is not supported.
- Ensure the presence of a home directory corresponding to the geo-rep user in every destination node.
- A user with the same user name should always exist under the home directory. If user was created using useradd command through CLI, then that user will automatically get synchronized in Red Hat Gluster Storage Console.

For more information on Geo-replication see the *Geo-replication* section in *Red Hat Gluster Storage Administration Guide*.

CHAPTER 11. USERS

This section describes the users in Red Hat Gluster Storage Console, how to set up user roles that control user permission levels, and how to manage users on the Red Hat Gluster Storage. Red Hat Gluster Storage Console relies on directory services for user authentication and information.

Users are assigned roles that allow them to perform their tasks as required. The role with the highest level of permissions is the admin role, which allows a user to set up, manage, and optimize all aspects of the Red Hat Gluster Storage Console. By setting up and configuring roles with permissions to perform actions and create objects, users can be provided with a range of permissions that allow the safe delegation of some administrative tasks to users without granting them complete administrative control.

Red Hat Gluster Storage Console provides a rich user interface that allows an administrator to manage their storage infrastructure from a web browser allowing even the most advanced configurations such as network bonding and VLANs to be centrally managed from a graphical console.



NOTE

Users are not created in Red Hat Gluster Storage Console, but in the Directory Services domain. Red Hat Gluster Storage Console can be configured to use multiple Directory Services domains.

11.1. DIRECTORY SERVICES SUPPORT IN RED HAT GLUSTER STORAGE CONSOLE

During installation, Red Hat Gluster Storage Console creates its own internal administration user, **admin**. This account is intended for use when initially configuring the environment, and for troubleshooting. To add other users to Red Hat Gluster Storage Console you will need to attach a directory server to the Console using the Domain Management Tool, **rhsc-manage-domains**.

Once at least one directory server has been attached to the Console you will be able to add users that exist in the directory server and assign roles to them using the Administration Portal. Users will be identified by their User Principle Name (UPN) of the form **user@domain**. Attachment of more than one directory server to the Console is also supported.

The directory servers currently supported for use with Red Hat Gluster Storage Console are:

- Active Directory;
- Identity Management (IdM); and
- Red Hat Directory Server(RHDS).


You must ensure that the correct DNS records exist for your directory server. In particular you must ensure that the DNS records for the directory server include:

- A valid pointer record (PTR) for the directory server's reverse look-up address.
- A valid service record (SRV) for LDAP over TCP port **389**.
- A valid service record (SRV) for Kerberos over TCP port **88**.
- A valid service record (SRV) for Kerberos over UDP port **88**.

If these records do not exist in DNS then you will be unable to add the domain to the Red Hat Gluster Storage Console configuration using **rhsc-manage-domains**.

For more detailed information on installing and configuring a supported directory server, refer to the vendor's documentation:

- Active Directory - <http://technet.microsoft.com/en-us/windowsserver/dd448614>.
- Identity Management (IdM) - https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/Identity_Management_Guide/index.html
- Red Hat Directory Server (RHDS) Documentation - https://access.redhat.com/site/documentation/en-US/Red_Hat_Directory_Server/



IMPORTANT

A user must be created in the directory server specifically for use as the Red Hat Gluster Storage administrative user. Do *not* use the administrative user for the directory server as the Red Hat Gluster Storage administrative user.



IMPORTANT

It is not possible to install Red Hat Gluster Storage Console (RHGSC) and IdM (ipa-server) on the same system. IdM is incompatible with the mod_ssl package, which is required by Red Hat Gluster Storage Console.

For information on creation of user accounts in Active Directory refer to <http://technet.microsoft.com/en-us/library/cc732336.aspx>.

For information on delegation of control in Active Directory refer to <http://technet.microsoft.com/en-us/library/cc732524.aspx>.

NOTE

Red Hat Gluster Storage Console uses Kerberos to authenticate with directory servers. RHDS does not provide native support for Kerberos. If you are using RHDS as your directory server then you must ensure that the directory server is made a service within a valid Kerberos domain. To do this you will need to perform these steps while referring to the relevant directory server documentation:

- Configure the **memberOf** plug-in for RHDS to allow group membership. In particular ensure that the value of the *memberofgroupattr* attribute of the **memberOf** plug-in is set to **uniqueMember**.

Consult the Red Hat Directory Server *Plug-in Guide* for more information on configuring the **memberOf** plug-in.

- Define the directory server as a service of the form **ldap/hostname@REALMNAME** in the Kerberos realm. Replace *hostname* with the fully qualified domain name associated with the directory server and *REALMNAME* with the fully qualified Kerberos realm name. The Kerberos realm name must be specified in capital letters.
- Generate a **keytab** file for the directory server in the Kerberos realm. The **keytab** file contains pairs of Kerberos principals and their associated encrypted keys. These keys will allow the directory server to authenticate itself with the Kerberos realm.

Consult the documentation for your Kerberos principle for more information on generating a **keytab** file.

- Install the **keytab** file on the directory server. Then configure RHDS to recognize the **keytab** file and accept Kerberos authentication using GSSAPI.

Consult the Red Hat Directory Server *Administration Guide* for more information on configuring RHDS to use an external **keytab** file.

- Test the configuration on the directory server by using the **kinit** command to authenticate as a user defined in the Kerberos realm. Once authenticated run the **ldapsearch** command against the directory server. Use the **-Y GSSAPI** parameters to ensure the use of Kerberos for authentication.

11.2. AUTHORIZATION MODEL

Red Hat Gluster Storage Console applies authorization controls to each action performed in the system. Authorization is applied based on the combination of the three components in any action:

- The user performing the action
- The type of action being performed
- The object on which the action is being performed

Actions

For an action to be successfully performed, the **user** must have the appropriate **permission** for the **object** being acted upon. Each type of action corresponds to a **permission**. There are many different permissions in the system, so for simplicity they are grouped together in **roles**.




Figure 11.1. Actions

Permissions

Permissions enable users to perform actions on objects, where objects are either individual objects or container objects.

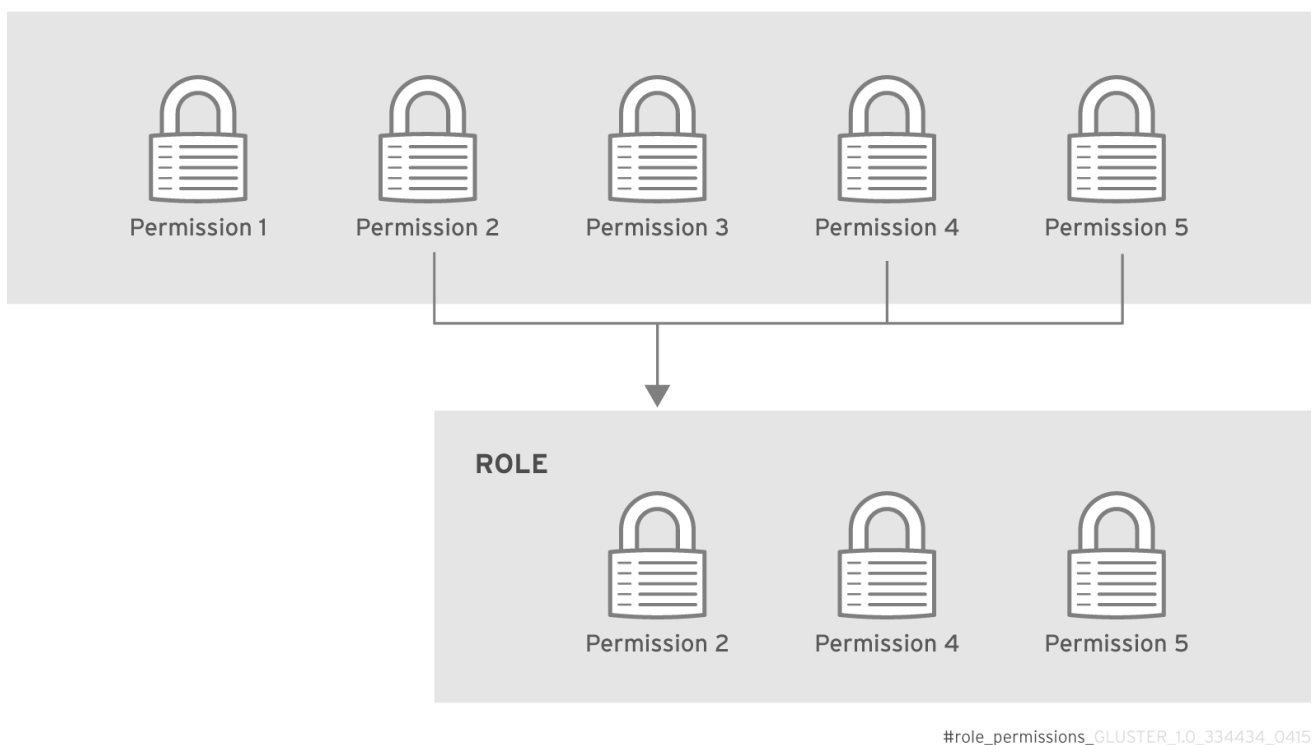



Figure 11.2. Permissions & Roles

Any permissions that apply to a container object also apply to all members of that container.



IMPORTANT

Some actions are performed on more than one object.

11.3. USER PROPERTIES

Roles and Permissions can be considered as the properties of the User object. Roles are predefined sets of privileges that can be configured from Red Hat Gluster Storage Console, permitting access and management to different levels of resources in the cluster, to specific physical and virtual resources. Multilevel administration includes a hierarchy of permissions that can be configured to provide a finely grained model of permissions, or a wider level of permissions as required by your enterprise. For

example, a cluster administrator has permissions to manage all servers in the cluster, while a server administrator has system administrator permissions to a single server. A user can have permissions to log into and use a single server but not make any changes to the server configurations, while another user can be assigned system permissions to a server, effectively acting as system administrator on the server.

11.3.1. Roles

Red Hat Gluster Storage provides a range of pre-configured or default roles, from the Superuser or system administration, to an end user with permissions to access a single volume only. There are two types of system administration roles: roles with system permissions to physical resources, such as hosts and storage; and roles with system permissions to virtual resources such as volumes. While you cannot change the default roles, you can clone them, and then customize the new roles as required.

Red Hat Gluster Storage Console has an **administrator** role. The privileges provided by this role are shown in this section.



NOTE

The default roles cannot be removed from the Red Hat Gluster Storage, or privileges cannot be modified; however the name and descriptions can be changed.

Administrator Role

- Allows access to the *Administration Portal* for managing servers and volumes.

For example, if a user has an **administrator** role on a cluster, they could manage all servers in the cluster using the *Administration Portal*.

Table 11.1. Red Hat Gluster Storage Console System Administrator Roles

Role	Privileges	Notes
SuperUser	Full permissions across all objects and levels	Can manage all objects across all clusters.
ClusterAdmin	Cluster Administrator	Can use, create, delete, and manage all resources in a specific cluster, including servers and volumes.
GlusterAdmin	Gluster Administrator	Can create, delete, configure and manage a specific volume. Can also add or remove host.
HostAdmin	Host Administrator	Can configure, manage, and remove a specific host. Can also perform network-related operations on a specific host.
NetworkAdmin	Network Administrator	Can configure and manage networks attached to servers.

11.3.2. Permissions

The following table details the actions for each object in the cluster, for each of which permission may be assigned. This results in a high level of control over actions at multiple levels.

Table 11.2. Permissions Actions on Objects

Object	Action
System - Configure RHS-C	Manipulate Users, Manipulate Permissions, Manipulate Roles, Generic Configuration
Cluster - Configure Cluster	Create, Delete, Edit Cluster Properties, Edit Network
Server - Configure Server	Create, Delete, Edit Host Properties, Manipulate Status, Edit Network
Gluster Storage - Configure Gluster Storage	Create, Delete, Edit Volumes, Volume Options, Manipulate Status

11.4. ASSIGNING AN ADMINISTRATOR OR USER ROLE TO A RESOURCE

Summary

Assign administrator or user roles to resources to allow users to access or manage that resource.

Procedure 11.1. Assigning a Role to a Resource

1. Click **Networks** tab and select a network from the results list.
2. Click **Permissions** sub-tab of the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Click **Add** to open the **Add Permission to User** window.

First Name	Last Name	User Name
------------	-----------	-----------

Figure 11.3. Add Permission to User

4. Enter the name or user name of an existing user into the **Search** text box and click **Go**. Select a user from the resulting list of possible matches.
5. Select a role from the **Role to Assign:** drop-down menu.
6. Click **OK** to assign the role and close the window.

Result

You have assigned a role to a user; the user now has the inherited permissions of that role enabled for that resource.

11.5. REMOVING AN ADMINISTRATOR OR USER ROLE FROM A RESOURCE

Summary

Remove an administrator or user role from a resource; the user loses the inherited permissions associated with the role for that resource.

Procedure 11.2. Removing a Role from a Resource

1. Click **Networks** tab and select a network from the results list.
2. Click **Permissions** sub-tab of the details pane to list the assigned users, the user's role, and the inherited permissions for the selected resource.
3. Select the user to remove from the resource.
4. Click **Remove**. The **Remove Permission** window opens to confirm permissions removal.
5. Click **OK** to remove the user role.

Result

You have removed the user's role, and the associated permissions, from the resource.

11.6. USERS OPERATIONS

Users can be added or removed from the system, assigned roles, and given permissions to various objects, enabling them to effectively perform their required work. The **Users** Details pane displays information on the status and privileges of users, enabling the system administrator to assign or change roles, allot servers, set up event notifications and allocate Directory Service groups. Because of the level of detail that is possible, a multi-level administration system can be defined.



NOTE

Login to the system is verified against the Directory Service records of the organization.

11.6.1. Adding Users and Groups

Existing users must be added to the Administration Portal before being assigned roles.

Adding Users

1. Click the **Users** tab. The list of authorized users for Red Hat Gluster Storage Console displays.
2. Click **Add**. The **Add Users and Groups** dialog box displays.

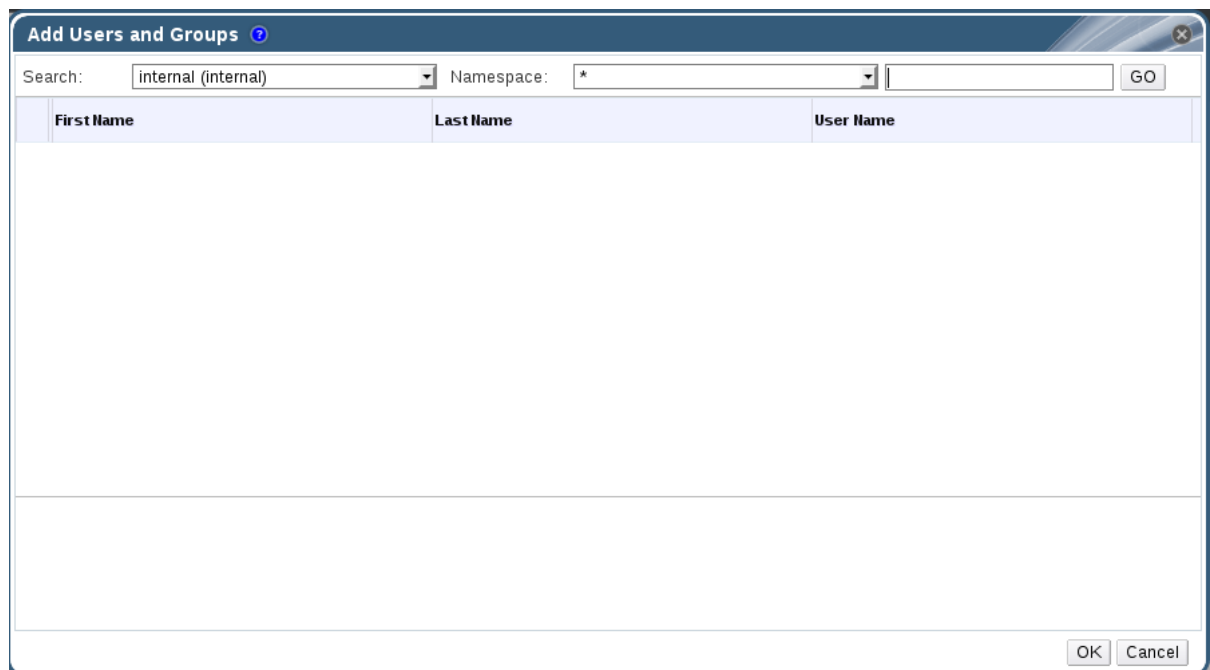


Figure 11.4. Add Users and Groups Dialog Box

3. The default **Search** domain displays. If there are multiple search domains, select the appropriate search domain. Enter a name or part of a name in the search text field, and click **GO**. Alternatively, click **GO** to view a list of all users and groups.
4. Select the group, user or users check boxes. The added user displays on the **Users** tab.

Viewing User Information

Users are not created from within the Red Hat Gluster Storage; Red Hat Gluster Storage Console

accesses user information from the organization's Directory Service. This means that you can only assign roles to users who already exist in your Directory Services domain. To assign permissions to users, use the **Permissions** tab on the Details pane of the relevant resource.

Example 11.1. Assigning a user permissions to use a particular server

To assign a user to a particular server, use the **Permissions** tab on the Details pane of the selected server.

To view general user information:

1. Click the **Users** tab. The list of authorized users for Red Hat Gluster Storage Console displays.
2. Select the user, or perform a search if the user is not visible on the results list.
3. The Details pane displays for the selected user, usually with the **General** tab displaying general information, such as the domain name, email, and status of the user.
4. The other tabs allow you to view groups, permissions, and events for the user.

For example, to view the groups to which the user belongs, click the **Directory Groups** tab.

11.6.2. Removing Users

A system administrator will need to remove users, for example, when they leave the company.

To remove a user:

1. Click the **Users** tab. The list of authorized users for Red Hat Gluster Storage Console displays.

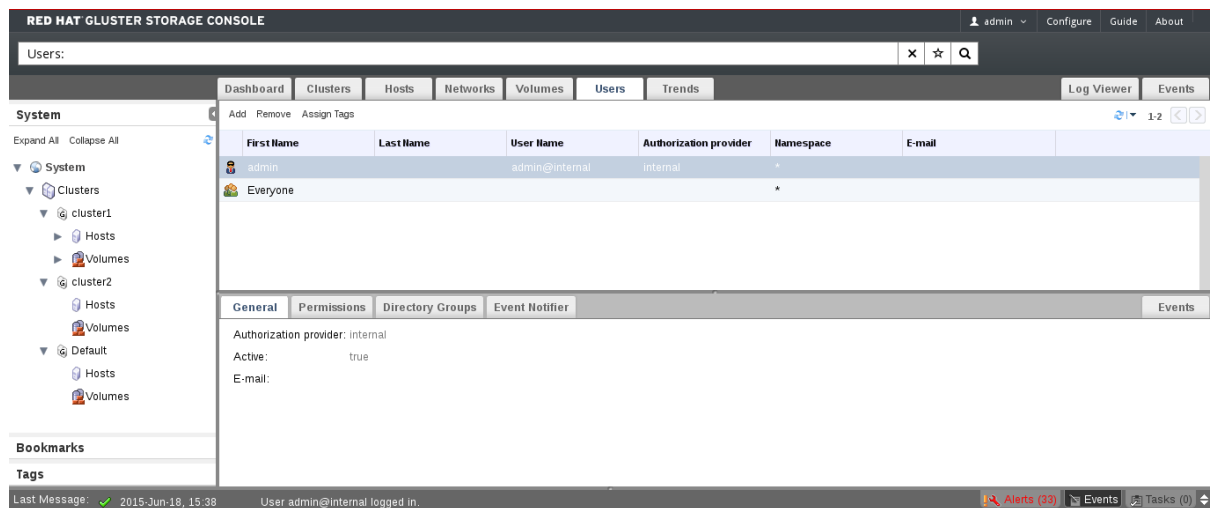


Figure 11.5. Users Tab

2. Select the user to be removed.
3. Click the **Remove** button. A message displays prompting you to confirm the removal.
4. Click **OK**.
5. The user is removed from Red Hat Gluster Storage Console.

**NOTE**

All user information is read from the Directory Service. Removing a user from the Red Hat Gluster Storage Console system deletes the record in the Red Hat Gluster Storage Console database, denying the user the ability to log on to the console. It removes the association in the Directory Service between the console and the user. All other user properties remain intact.

11.7. EVENT NOTIFICATIONS

When troubleshooting problems related to users, the first thing to recall is that users must be correctly added to and authenticated at the Directory Services level, not on the Red Hat Gluster Storage Console. Problems with permissions can occur when adequate levels of permissions to all required objects have not been assigned. Users, particularly those with administrator roles, require to be notified when events or triggers occur.

11.7.1. Managing Event Notifiers

This section describes how to set up and manage event notifications for users. Events are displayed on the **Events** tab, however, users can be notified by email about selected events. For example, a system administrator might like to know when there is a problem with storage, or a team lead may want to be notified if a volume is down.

To set up event notifications:

1. Click the **Users** tab. The list of authorized users for Red Hat Gluster Storage Console displays.
2. Select the user who requires notification, or perform a search if the user is not visible on the results list.
3. Click the **Event Notifier** tab. The **Event Notifier** tab displays a list of events for which the user will be notified, if any.
4. Click the **Manage Events** button. The **Add Event Notification** dialog box displays a list of events for Services, Hosts, Volumes, Hooks, and General Management events. You can select all, or pick individual events from the list. Click the **Expand All** button to see complete lists of events.

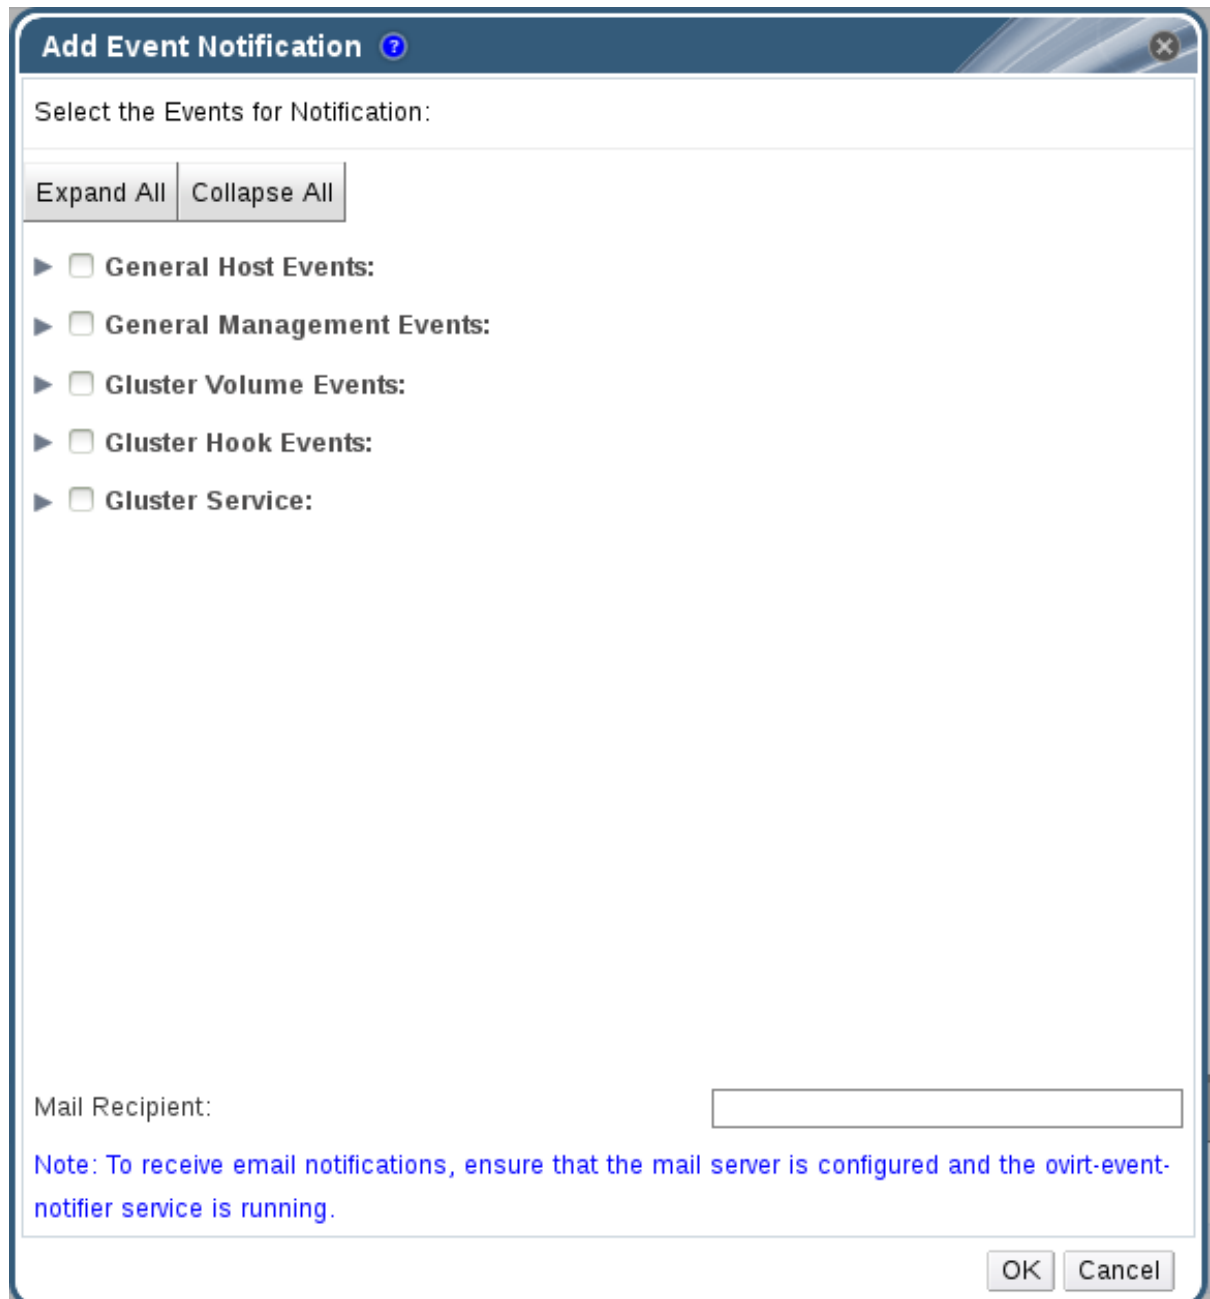



Figure 11.6. The Add Events Dialog Box

5. Enter an email address in the **Mail Recipient:** field.
6. Click **OK** to save changes and close the window. The selected events display on the **Event Notifier** tab for the user.
7. Configure the **ovirt-engine-notifier** service on the Red Hat Gluster Storage Console.



IMPORTANT

The *MAIL_SERVER* parameter is mandatory.

The event notifier configuration file can be found in `/usr/share/ovirt-engine/services/ovirt-engine-notifier/ovirt-engine-notifier.conf`. The parameters for event notifications in *ovirt-engine-notifier.conf* are listed in [Table 11.3, “ovirt-engine-notifier.conf variables”](#)

Table 11.3. ovirt-engine-notifier.conf variables

Variable name	Default	Remarks
INTERVAL_IN_SECONDS	120	The interval in seconds between instances of dispatching messages to subscribers.
MAIL_SERVER	none	The SMTP mail server address. Required.
MAIL_PORT	25	The default port of a non-secured SMTP server is 25. The default port of a secured SMTP server (one with SSL enabled) is 465.
MAIL_USER	none	If SSL is enabled to authenticate the user, then this variable must be set. This variable is also used to specify the "from" user address when the MAIL_FROM variable is not set. Some mail servers do not support this functionality. The address is in RFC822 format.
MAIL_PASSWORD	none	This variable is required to authenticate the user if the mail server requires authentication or if SSL is enabled.
MAIL_ENABLE_SSL	false	This indicates whether SSL should be used to communicate with the mail server.
HTML_MESSAGE_FORMAT	false	The mail server sends messages in HTML format if this variable is set to "true".
MAIL_FROM	none	This variable specifies a "from" address in RFC822 format, if supported by the mail server.
MAIL_REPLY_TO	none	This variable specifies "reply-to" addresses in RFC822 format on sent mail, if supported by the mail server.
DAYS_TO_KEEP_HISTORY	none	This variable sets the number of days dispatched events will be preserved in the history table. If this variable is not set, events remain on the history table indefinitely.

Variable name	Default	Remarks
DAYS_TO_SEND_ON_STARTUP	0	This variable specifies the number of days of old events that are processed and sent when the notifier starts. If set to 2, for example, the notifier will process and send the events of the last two days. Older events will just be marked as processed and won't be sent. The default is 0, so no old messages will be sent at all during startup.

8. Start the **ovirt-engine-notifier** service on the Red Hat Gluster Storage Console. This activates the changes you have made:

```
# /etc/init.d/ovirt-engine-notifier start
```

You now receive emails based on events in your Red Hat Gluster Storage Environment. The selected events display on the Event Notifier tab for the user.

To cancel event notification:

1. In the **Users** tab, select the user or the user group.
2. Select the **Event Notifier** tab. The Details pane displays the events for which the user will receive notifications.
3. Click the **Manage Events** button. The **Add Event Notification** dialog box displays a list of events for Servers, Gluster Volume, and General Management events. To remove an event notification, deselect events from the list. Click the **Expand All** button to see the complete lists of events.
4. Click **OK**. The deselected events are removed from the display on the **Event Notifier** tab for the user.

PART III. MONITORING

CHAPTER 12. MONITORING RED HAT GLUSTER STORAGE CONSOLE

System administrators monitor the management environment to view the overall performance of cluster infrastructure. This helps identify key areas in the cluster environment that require attention or optimization. System administrators gain up-to-date information on the performance and status of storage environment components with the **Events** list.

The **Events** list lists all warnings, errors, and other events that occur in the system.

12.1. VIEWING THE EVENT LIST

The **Events** list displays all system events. You can view the events by clicking on **Events** tab. The type of events that appear in the Events tab are audits, warnings, and errors. The names of the user, host, cluster, and Gluster volume involved in the event are also listed. This information helps determine the cause of the event. Click column headers to sort the event list.

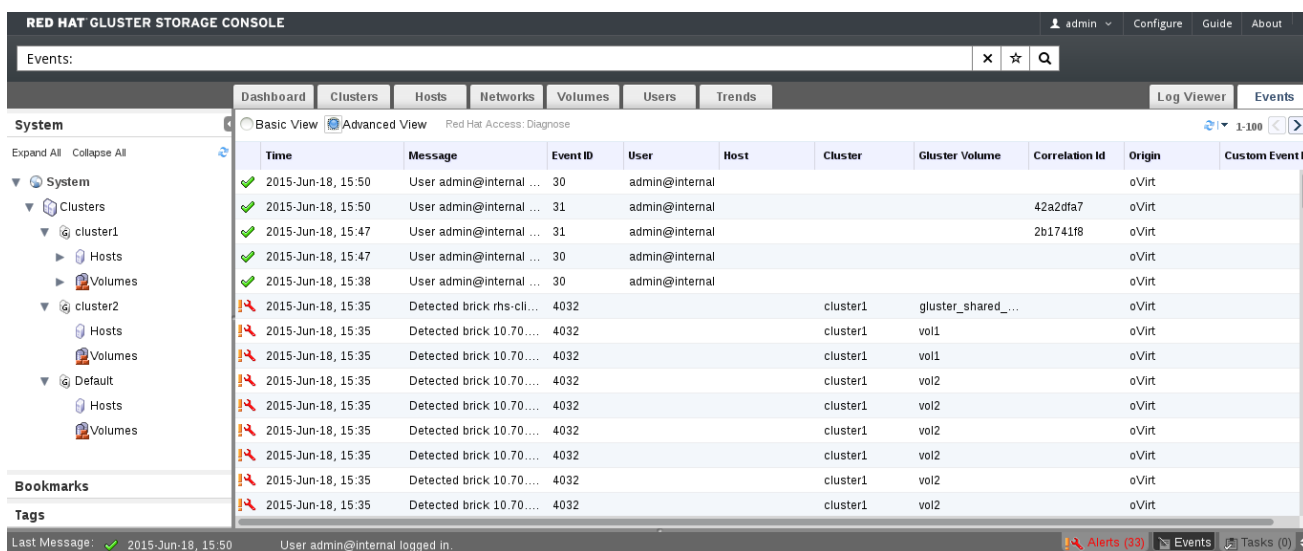


Figure 12.1. Event List - Advanced View

The following table describes the different columns of the **Event** list:

Column	Description
Event	The type of event. The possible event types are: <ul style="list-style-type: none"> Audit notification (e.g. log on). Warning notification. Error notification.
Time	The time that the event occurred.
Message	The message describing that an event occurred.

Column	Description
User	The user that received the event.
Host	The host on which the event occurred.
Cluster	The cluster on which the event occurred.

12.2. VIEWING ALERT INFORMATION

The **Alerts** pane lists all important notifications regarding the cluster environment.

Alerts display in the lowermost panel on top of the manager interface. Drag the top of the **Alert** pane to resize it or click the minimize/maximize icon in the top right of the pane to show or hide it.

The **Alerts** pane also contains an **Events** list. Click the **Events** tab in the **Alerts** pane to display the **Events** list. See [Section 12.1, “Viewing the Event List”](#) for more information about the **Events** list.

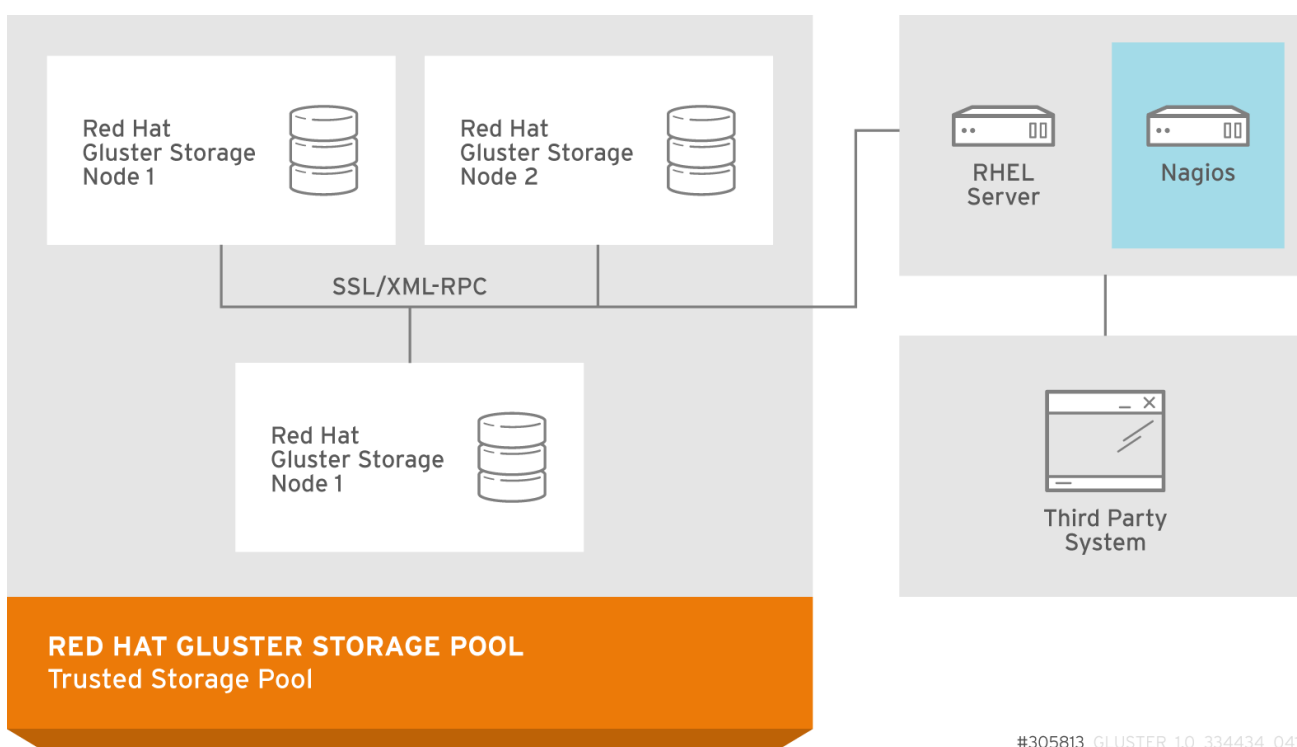
CHAPTER 13. MONITORING RED HAT GLUSTER STORAGE USING NAGIOS

Red Hat Gluster Storage Console provides monitoring of Red Hat Gluster Storage trusted storage pool by building on the Nagios platform. You can view the physical and logical resource utilization and status (CPU, Memory, Disk, Network, Swap, Cluster, Volume, Brick) in Trends tab of Red Hat Gluster Storage Console. Nagios is installed and enabled in Red Hat Gluster Storage Console Server by default to monitor Red Hat Gluster Storage nodes. To monitor using Nagios, add hosts to Red Hat Gluster Storage Console and configure Nagios using auto-discovery. For more information about adding hosts, see [Section 6.2.1, “Adding Hosts”](#). This chapter describes the procedures for deploying Nagios on Red Hat Gluster Storage Console node

For more information on Nagios, see *Nagios Documentation*.

For more information on Changing Nagios Password and Creating Nagios User, see the corresponding sections in *Red Hat Gluster Storage Administration Guide*.

The following diagram illustrates deployment of Nagios on Red Hat Gluster Storage Console Server.



#305813_GLUSTER_1.0_334434_0415

Figure 13.1. Nagios deployed on Red Hat Gluster Storage Console Server

13.1. CONFIGURING NAGIOS

Auto-Discovery is a python script which automatically discovers all the nodes and volumes in the cluster. It also creates Nagios configuration to monitor them. By default, it runs once in 24 hours to synchronize the Nagios configuration from Red Hat Gluster Storage Trusted Storage Pool configuration.

For more information on Nagios Configuration files, see [Section D.1, “Nagios Configuration Files”](#)



NOTE

Before configuring Nagios using `configure-gluster-nagios` command, ensure that all the Red Hat Gluster Storage nodes are configured.

1. Execute **configure-gluster-nagios** command manually only the first time with cluster name and host address on the Nagios server using the following command:

```
# configure-gluster-nagios -c cluster-name -H HostName-or-IP-address
```

For **-c**, provide a cluster name (a logical name for the cluster) and for **-H**, provide the host name or ip address of a node in the Red Hat Gluster Storage trusted storage pool.

2. Perform the steps given below when **configure-gluster-nagios** command runs:
 1. Confirm the configuration when prompted.
 2. Enter the current Nagios server host name or IP address to be configured all the nodes.
 3. Confirm restarting Nagios server when prompted.

```
# configure-gluster-nagios -c demo-cluster -H HostName-or-IP-address
Cluster configurations changed
Changes :
Hostgroup demo-cluster - ADD
Host demo-cluster - ADD
Service - Volume Utilization - vol-1 -ADD
Service - Volume Split-Brain - vol-1 -ADD
Service - Volume Status - vol-1 -ADD
Service - Volume Utilization - vol-2 -ADD
Service - Volume Status - vol-2 -ADD
Service - Cluster Utilization -ADD
Service - Cluster - Quorum -ADD
Service - Cluster Auto Config -ADD
Host Host_Name - ADD
Service - Brick Utilization - /bricks/vol-1-5 -ADD
Service - Brick - /bricks/vol-1-5 -ADD
Service - Brick Utilization - /bricks/vol-1-6 -ADD
Service - Brick - /bricks/vol-1-6 -ADD
Service - Brick Utilization - /bricks/vol-2-3 -ADD
Service - Brick - /bricks/vol-2-3 -ADD
Are you sure, you want to commit the changes? (Yes, No) [Yes]:
Enter Nagios server address [Nagios_Server_Address]:
Cluster configurations synced successfully from host ip-address
Do you want to restart Nagios to start monitoring newly
discovered entities? (Yes, No) [Yes]:
Nagios re-started successfully
```

All the hosts, volumes and bricks are added and displayed.

3. Login to the Nagios server GUI using the following URL.

```
https://NagiosServer-HostName-or-IPaddress/nagios
```



NOTE

- The default Nagios user name and password is *nagiosadmin / nagiosadmin*.
- You can manually update/discover the services by executing the **configure-gluster-nagios** command or by running **Cluster Auto Config** service through Nagios Server GUI.
- If the node with which auto-discovery was performed is down or removed from the cluster, run the **configure-gluster-nagios** command with a different node address to continue discovering or monitoring the nodes and services.
- If new nodes or services are added, removed, or if snapshot restore was performed on Red Hat Gluster Storage node, run **configure-gluster-nagios** command.

13.2. CONFIGURING NAGIOS SERVER FOR SEND MAIL NOTIFICATIONS

1. In the `/etc/nagios/gluster/gluster-contacts.cfg` file, add contacts to send mail in the format shown below:

Modify **contact_name**, **alias**, and **email**.

```

define contact {
    contact_name           Contact1
    alias                  ContactNameAlias
    email                  email-address
    service_notification_period 24x7
    service_notification_options w,u,c,r,f,s
    service_notification_commands notify-service-by-
email
    host_notification_period 24x7
    host_notification_options d,u,r,f,s
    host_notification_commands notify-host-by-
email
}
define contact {
    contact_name           Contact2
    alias                  ContactNameAlias2
    email                  email-address
    service_notification_period 24x7
    service_notification_options w,u,c,r,f,s
    service_notification_commands notify-service-by-
email
    host_notification_period 24x7
    host_notification_options d,u,r,f,s
    host_notification_commands notify-host-by-
email
}


```

The **service_notification_options** directive is used to define the service states for which notifications can be sent out to this contact. Valid options are a combination of one or more of the following:

- **w**: Notify on WARNING service states
- **u**: Notify on UNKNOWN service states
- **c**: Notify on CRITICAL service states
- **r**: Notify on service RECOVERY (OK states)
- **f**: Notify when the service starts and stops FLAPPING
- **n (none)**: Do not notify the contact on any type of service notifications

The **host_notification_options** directive is used to define the host states for which notifications can be sent out to this contact. Valid options are a combination of one or more of the following:

- **d**: Notify on DOWN host states
- **u**: Notify on UNREACHABLE host states
- **r**: Notify on host RECOVERY (UP states)
- **f**: Notify when the host starts and stops FLAPPING
- **s**: Send notifications when host or service scheduled downtime starts and ends
- **n (none)**: Do not notify the contact on any type of host notifications.



NOTE

By default, a contact and a contact group are defined for administrators in **contacts.cfg** and all the services and hosts will notify the administrators. Add suitable email id for administrator in **contacts.cfg** file.

2. To add a group to which the mail need to be sent, add the details as given below:

```

define contactgroup{
    contactgroup_name      Group1
    alias                  GroupAlias
    members                Contact1, Contact2
}

```

3. In the **/etc/nagios/gluster/gluster-templates.cfg** file specify the contact name and contact group name for the services for which the notification need to be sent, as shown below:

Add **contact_groups** name and **contacts** name.

```

define host{
    name                gluster-generic-host
    use                  linux-server
    notifications_enabled 1
}

```

```

notification_period      24x7
notification_interval    120
notification_options     d,u,r,f,s
register                 0
contact_groups           Group1
contacts                 Contact1,Contact2
}

define service {
name                    gluster-service
use                    generic-service
notifications_enabled   1
notification_period     24x7
notification_options     w,u,c,r,f,s
notification_interval   120
register                0
_gluster_entity        Service
contact_groups          Group1
contacts                Contact1,Contact2
}

```

You can configure notification for individual services by editing the corresponding node configuration file. For example, to configure notification for brick service, edit the corresponding node configuration file as shown below:

```

define service {
use                    brick-service
_VOL_NAME             VolumeName
__GENERATED_BY_AUTOCONFIG 1
notes                 Volume : VolumeName
host_name              RedHatStorageNodeName
_BRICK_DIR            brickpath
service_description    Brick Utilization - brickpath
contact_groups         Group1
contacts               Contact1,Contact2
}

```

- To receive detailed information on every update when Cluster Auto-Config is run, edit `/etc/nagios/objects/commands.cfg` file add `$NOTIFICATIONCOMMENT$\n` after `$SERVICEOUTPUT$\n` option in `notify-service-by-email` and `notify-host-by-email` command definition as shown below:

```

# 'notify-service-by-email' command definition
define command{
command_name    notify-service-by-email
command_line    /usr/bin/printf "%b" "***** Nagios
*****\n\nNotification Type: $NOTIFICATIONTYPE$\n\nService:
$SERVICEDESC$\nHost: $HOSTALIAS$\nAddress: $HOSTADDRESS$\nState:
$SERVICESTATE$\n\nDate/Time: $LONGDATETIME$\n\nAdditional
Info:\n\n$SERVICEOUTPUT$\n $NOTIFICATIONCOMMENT$\n" | /bin/mail -s
"*** $NOTIFICATIONTYPE$ Service Alert: $HOSTALIAS$/$SERVICEDESC$ is
$SERVICESTATE$ ***" $CONTACTEMAIL$
}

```


This will send emails similar to the following when the service alert is triggered.

```
** PROBLEM Service Alert: vm912_6313/Volume Heal info - dist-rep3 is
WARNING **
***** Nagios *****

Notification Type: PROBLEM

Service: Volume Heal info - dist-rep3
Host: vm912_6313 (--for services, host is the cluster name)
Address: vm912_6313
State: WARNING

Date/Time: Sun May 15 15:03:39 IST 2016

Additional Info:

Unsynced entries found.
```

- Restart the Nagios server using the following command:

```
#service nagios restart
```

The Nagios server sends notifications during status changes to the mail addresses specified in the file.



NOTE

- By default, the system ensures three occurrences of the event before sending mail notifications.
- By default, Nagios Mail notification is sent using `/bin/mail` command. To change this, modify the definition for **notify-host-by-email** command and **notify-service-by-email** command in `/etc/nagios/objects/commands.cfg` file and configure the mail server accordingly.

13.3. VERIFYING THE CONFIGURATION

- Verify the updated configurations using the following command:

```
# nagios -v /etc/nagios/nagios.cfg
```


If error occurs, verify the parameters set in `/etc/nagios/nagios.cfg` and update the configuration files.

- Restart Nagios server using the following command:

```
# service nagios restart
```

- Log into the Nagios server GUI using the following URL with the Nagios Administrator user name and password.

```
https://NagiosServer-HostName-or-IPaddress/nagios
```



NOTE

To change the default password, see *Changing Nagios Password* section in *Red Hat Gluster Storage Administration Guide*.

4. Click **Services** in the left pane of the Nagios server GUI and verify the list of hosts and services displayed.

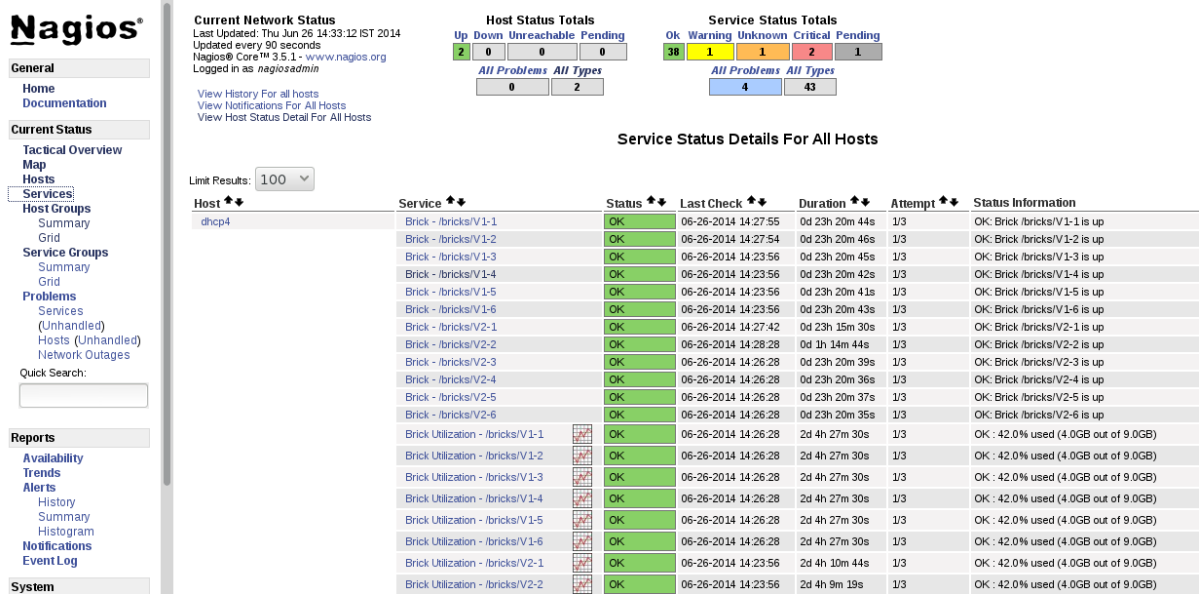


Figure 13.2. Nagios Services

13.4. USING NAGIOS SERVER GUI

You can monitor Red Hat Gluster Storage trusted storage pool through Nagios Server GUI.

To view the details, log into the Nagios Server GUI by using the following URL.

```
https://NagiosServer-HostName-or-IPAddress/nagios
```

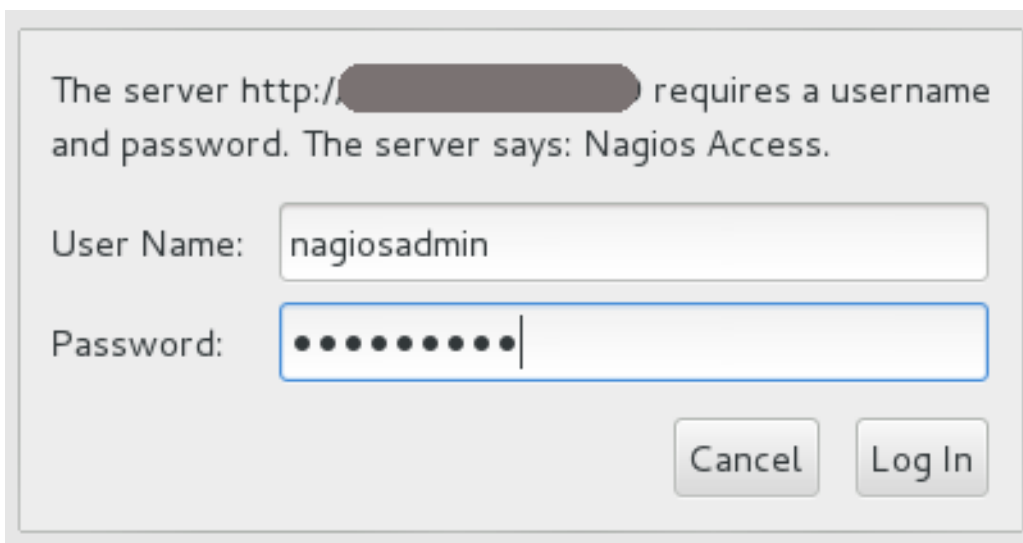


Figure 13.3. Nagios Login

Cluster Overview

To view the overview of the hosts and services being monitored, click **Tactical Overview** in the left pane. The overview of Network Outages, Hosts, Services, and Monitoring Features are displayed.

The screenshot shows the Nagios Tactical Monitoring Overview page. The left sidebar contains navigation links for General, Current Status (Tactical Overview, Map, Hosts, Services, Host Groups, Service Groups, Problems), and Reports. The main content area is titled 'Tactical Monitoring Overview' and includes a 'Monitoring Performance' box with execution and latency times. Below this are sections for 'Network Outages' (0 Outages), 'Hosts' (0 Down, 0 Unreachable, 2 Up, 0 Pending), and 'Services' (2 Critical, 1 Warning, 1 Unknown, 38 Ok, 1 Pending). A 'Monitoring Features' table shows the status of various features like Flap Detection, Notifications, Event Handlers, Active Checks, and Passive Checks.

Flap Detection	Notifications	Event Handlers	Active Checks	Passive Checks
✓ All Services Enabled	✓ All Services Enabled	✓ All Services Enabled	✓ 1 Service Disabled	✓ All Services Enabled
No Services Flapping	All Hosts Enabled	All Hosts Enabled	All Hosts Enabled	All Hosts Enabled
2 Hosts Disabled	No Hosts Flapping			

Figure 13.4. Tactical Overview

Host Status

To view the status summary of all the hosts, click **Summary** under **Host Groups** in the left pane.

The screenshot shows the Nagios Host Groups Summary page. The left sidebar is similar to the previous page, but 'Summary' under 'Host Groups' is selected. The main content area shows 'Current Network Status' and 'Host Status Totals' (Up: 5, Down: 0, Unreachable: 0, Pending: 0). Below this is a 'Service Status Totals' summary (Ok: 47, Warning: 0, Unknown: 3, Critical: 1, Pending: 0). The central part of the page is titled 'Status Summary For All Host Groups' and contains a table with columns for Host Group, Host Status Summary, and Service Status Summary.

Host Group	Host Status Summary	Service Status Summary
cluster-1 (cluster-1)	3 UP	38 OK 3 UNKNOWN : 3 Unhandled
gluster_hosts (gluster_hosts)	3 UP	38 OK 3 UNKNOWN : 3 Unhandled

Figure 13.5. Host Groups Summary

To view the list of all hosts and their status, click **Hosts** in the left pane.



- General**
- Home
- Documentation
- Current Status**
- Tactical Overview
- Map
- Hosts
- Services
- Host Groups
- Summary
- Grid
- Service Groups
- Summary
- Grid
- Problems
- Services (Unhandled)
- Hosts (Unhandled)
- Network Outages
- Quick Search:
-
- Reports**
- Availability
- Trends
- Alerts
- History
- Summary

Current Network Status
 Last Updated: Thu Jun 26 17:04:31 IST 2014
 Updated every 90 seconds
 Nagios® Core™ 3.5.1 - www.nagios.org
 Logged in as nagiosadmin

Host Status Totals

Up	Down	Unreachable	Pending
3	0	0	0
All Problems All Types			
0		3	

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
40	0	1	2	1
All Problems All Types				
3			44	

View Service Status Detail For All Host Groups
 View Status Overview For All Host Groups
 View Status Summary For All Host Groups
 View Status Grid For All Host Groups

Host Status Details For All Host Groups

Limit Results:

Host	Status	Last Check	Duration	Status Information
	UP	06-26-2014 17:01:03	0d 0h 6m 19s	OK: Host is UP
temp_node1	UP	06-26-2014 17:00:32	0d 0h 12m 16s	PING OK - Packet loss = 0%, RTA = 0.03 ms
test-cluster	UP	06-26-2014 17:02:12	0d 0h 7m 18s	OK : None of the Volumes in the cluster are in Critical State

Results 1 - 3 of 3 Matching Hosts

Figure 13.6. Host Status

Service Status

To view the list of all hosts and their service status click **Services** in the left pane.

Host	Service	Status	Last Check	Duration	Attempt	Status Information	
30.70.47.40	CTDB	UNKNOWN	06-27-2014 00:31:51	2d 12h 5m 40s	3/3	CTDB not configured	
	Cpu Utilization	OK	06-27-2014 12:33:17	1d 12h 32m 0s	1/3	CPU Status OK: Total CPU 3.6% idle CPU 96.40%	
	Disk Utilization	OK	06-27-2014 12:33:51	2d 12h 4m 26s	1/3	OK: 30.0% used (3.0GB out of 10.0GB)	
	Cluster Management	OK	06-27-2014 00:35:58	2d 12h 4m 26s	1/3	Process glusterd is running	
	Memory Utilization	OK	06-27-2014 12:33:25	2d 12h 4m 26s	1/3	OK: 34.55% used(1.34GB out of 3.87GB)	
	NFS	OK	06-27-2014 00:40:52	2d 12h 4m 26s	1/3	Process glusterfs-ntfs is running	
	Network Utilization	OK	06-27-2014 12:33:19	2d 12h 4m 26s	1/3	OK: ovirtmgmt UP	
	Quota	OK	06-27-2014 00:45:46	2d 12h 4m 26s	1/3	Process quotad is running	
	SMB	OK	06-27-2014 00:48:13	2d 12h 4m 26s	1/3	Process smb is running	
	Self-Heal	OK	06-27-2014 00:50:40	2d 12h 4m 26s	1/3	Gluster Self-Heal Daemon is running	
	Swap Utilization	OK	06-27-2014 12:33:17	2d 12h 4m 26s	1/3	OK: 0.00% used(0.00GB out of 1.00GB)	
	Brick - idisk1vol1-a	OK	06-27-2014 12:26:26	2d 12h 7m 51s	1/3	OK: Brick idisk1vol1-a is up	
	Brick - /exportvol2-a	OK	06-27-2014 12:28:01	2d 12h 7m 51s	1/3	OK: Brick /exportvol2-a is up	
	Brick Utilization - idisk1vol1-a	OK	06-27-2014 12:29:14	2d 12h 7m 51s	1/3	OK: 38.0% used (1.0GB out of 2.6GB)	
Brick Utilization - /exportvol2-a	OK	06-27-2014 12:31:41	2d 12h 7m 51s	1/3	OK: 42.0% used (4.0GB out of 9.0GB)		
30.70.47.41	CTDB	UNKNOWN	06-27-2014 00:34:08	2d 12h 7m 51s	3/3	CTDB not configured	
	Cpu Utilization	OK	06-27-2014 12:33:27	0d 12h 31m 51s	1/3	CPU Status OK: Total CPU 13.55% idle CPU 86.44%	
	Disk Utilization	OK	06-27-2014 12:33:27	2d 12h 7m 51s	1/3	OK: 50.0% used (5.0GB out of 10.0GB)	
	Cluster Management	OK	06-27-2014 00:41:29	2d 12h 7m 51s	1/3	Process glusterd is running	
	Memory Utilization	OK	06-27-2014 12:33:27	0d 12h 31m 51s	1/3	OK: 73.46% used(2.84GB out of 3.87GB)	
	NFS	OK	06-27-2014 00:46:23	2d 12h 7m 51s	1/3	Process glusterfs-ntfs is running	
	Network Utilization	OK	06-27-2014 12:33:27	0d 12h 31m 51s	1/3	OK: ovirtmgmt UP	
	Quota	OK	06-27-2014 00:51:17	2d 12h 7m 51s	1/3	Process quotad is running	
	SMB	OK	06-27-2014 00:53:44	2d 12h 7m 51s	1/3	Process smb is running	
	Self-Heal	OK	06-27-2014 00:56:11	2d 12h 7m 51s	1/3	Gluster Self-Heal Daemon is running	
	Swap Utilization	OK	06-27-2014 12:33:27	0d 12h 31m 51s	1/3	OK: 0.00% used(0.00GB out of 1.00GB)	
	Brick - idisk1vol1-b	OK	06-27-2014 12:29:51	2d 12h 4m 48s	1/3	OK: Brick idisk1vol1-b is up	
	Brick - /exportvol2-b	OK	06-27-2014 12:32:18	2d 12h 4m 48s	1/3	OK: Brick /exportvol2-b is up	
	Brick Utilization - idisk1vol1-b	OK	06-27-2014 12:24:45	2d 12h 4m 48s	1/3	OK: 29.0% used (1.0GB out of 2.6GB)	
Brick Utilization - /exportvol2-b	OK	06-27-2014 12:29:29	2d 12h 4m 48s	1/3	OK: 35.0% used (3.0GB out of 9.0GB)		
30.70.47.42	CTDB	UNKNOWN	06-27-2014 00:38:36	2d 12h 4m 48s	3/3	CTDB not configured	
	Cpu Utilization	OK	06-27-2014 12:33:29	0d 12h 31m 48s	1/3	CPU Status OK: Total CPU 4.77% idle CPU 95.23%	
	Disk Utilization	OK	06-27-2014 12:33:29	2d 12h 4m 48s	1/3	OK: 40.0% used (4.0GB out of 10.0GB)	
	Cluster Management	OK	06-27-2014 00:46:59	2d 12h 4m 48s	1/3	Process glusterd is running	
	Memory Utilization	OK	06-27-2014 12:33:27	0d 12h 31m 51s	1/3	OK: 35.23% used(1.36GB out of 3.87GB)	
	NFS	OK	06-27-2014 00:51:53	2d 12h 4m 48s	1/3	Process glusterfs-ntfs is running	
	Network Utilization	OK	06-27-2014 12:33:29	0d 12h 31m 48s	1/3	OK: ovirtmgmt UP	
	Quota	OK	06-27-2014 00:56:47	2d 12h 4m 48s	1/3	Process quotad is running	
	SMB	OK	06-27-2014 00:59:14	2d 12h 4m 48s	1/3	Process smb is running	
	Self-Heal	OK	06-27-2014 00:30:28	2d 12h 4m 48s	1/3	Gluster Self-Heal Daemon is running	
	Swap Utilization	OK	06-27-2014 12:33:29	0d 12h 31m 48s	1/3	OK: 0.00% used(0.00GB out of 1.00GB)	
	cluster-1	Cluster - Quorum	?	06-25-2014 00:31:01	2d 12h 3m 27s	1/3	QUORUM: Server quorum regained for volume vol1. Storing local bricks.
		Cluster Auto Config	OK	06-27-2014 00:36:22	2d 12h 5m 24s	1/3	Cluster configurations are in sync
		Cluster Utilization	OK	06-27-2014 12:33:53	2d 12h 3m 24s	1/3	OK: used 35% of available 22.3098526001 GB
Volume Quota - vol1		CRITICAL	06-27-2014 12:32:53	0d 10h 13m 24s	3/3	QUOTA hard link reached on 1	
Volume Self-Heal - vol2		OK	06-27-2014 12:32:42	2d 12h 5m 24s	1/3	No unsynced entries present	
Volume Status - vol1		OK	06-27-2014 12:25:09	2d 12h 5m 24s	1/3	OK: Volume : DISTRIBUTE type - All bricks are Up	
Volume Status - vol2		OK	06-27-2014 12:28:53	2d 12h 5m 24s	1/3	OK: Volume : REPLICATE type - All bricks are Up	
Volume Utilization - vol1	OK	06-27-2014 12:31:32	2d 12h 5m 24s	1/3	OK: Utilization:30.88%		


Figure 13.7. Service Status



NOTE

In the left pane of Nagios Server GUI, click **Availability** and **Trends** under the **Reports** field to view the Host and Services Availability and Trends.

Host Services

1. Click **Hosts** in the left pane. The list of hosts are displayed.
2. Click  corresponding to the host name to view the host details.
3. Select the service name to view the Service State Information. You can view the utilization of the following services:
 - o Memory
 - o Swap
 - o CPU
 - o Network
 - o Brick
 - o Disk


The Brick/Disk Utilization Performance data has four sets of information for every mount point which are brick/disk space detail, inode detail of a brick/disk, thin pool utilization and thin pool metadata utilization if brick/disk is made up of thin LV.

The Performance data for services is displayed in the following format:
value[UnitOfMeasurement];warningthreshold;criticalthreshold;min;max.

For Example,

```
Performance Data: /bricks/brick2=31.596%;80;90;0;0.990
/bricks/brick2.inode=0.003%;80;90;0;1048064
/bricks/brick2.thinpool=19.500%;80;90;0;1.500 /bricks/brick2.thinpool-
metadata=4.100%;80;90;0;0.004
```

As part of disk utilization service, the following mount points will be monitored: / , **/boot** , **/home** , **/var** and **/usr** if available.

4. To view the utilization graph, click  corresponding to the service name. The utilization graph is displayed.

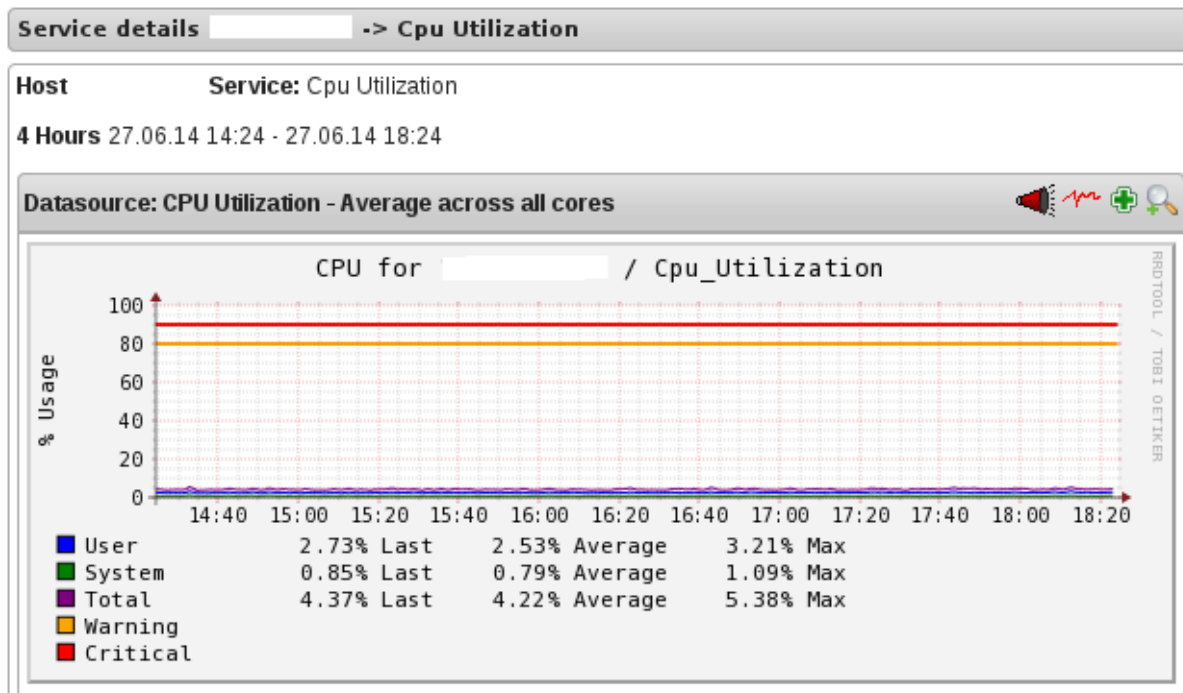


Figure 13.8. CPU Utilization

5. To monitor status, click on the service name. You can monitor the status for the following resources:
 - o Disk
 - o Network


6. To monitor process, click on the process name. You can monitor the following processes:
 - o Gluster NFS (Network File System)
 - o Self-Heal (Self Heal)
 - o Gluster Management (glusterd)
 - o Quota (Quota daemon)
 - o CTDB
 - o SMB




NOTE

Monitoring Openstack Swift operations is not supported.

Cluster Services

1. Click **Hosts** in the left pane. The list of hosts and clusters are displayed.
2. Click  corresponding to the cluster name to view the cluster details.

3. To view utilization graph, click  corresponding to the service name. You can monitor the following utilizations:
- o Cluster
 - o Volume

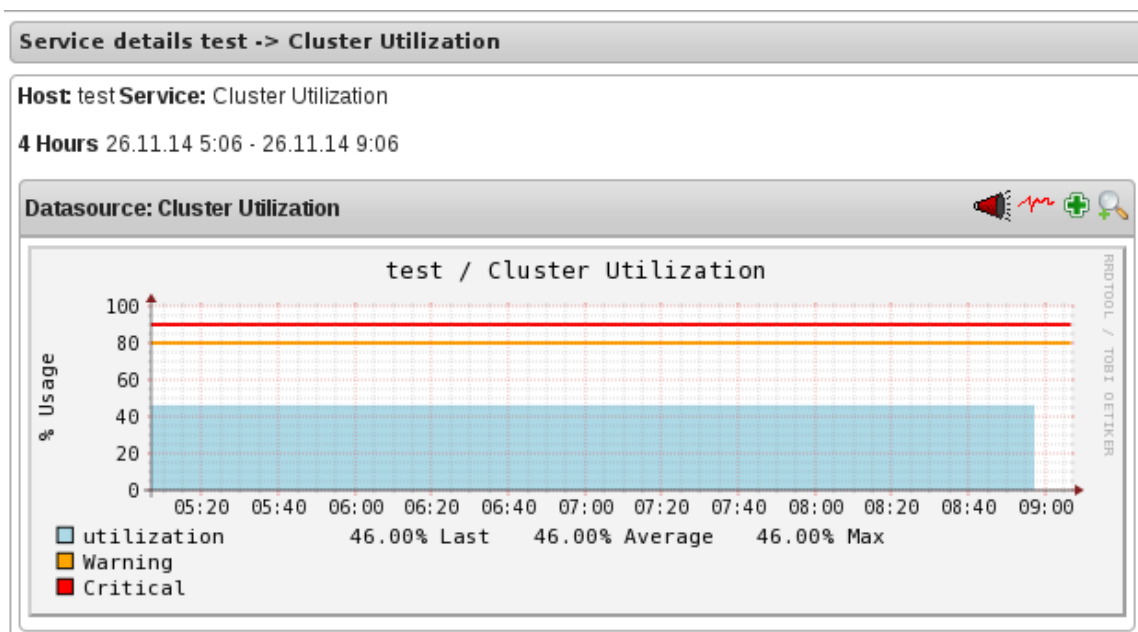


Figure 13.9. Cluster Utilization

4. To monitor status, click on the service name. You can monitor the status for the following resources:
- o Host
 - o Volume
 - o Brick
5. To monitor cluster services, click on the service name. You can monitor the following:
- o Volume Quota
 - o Volume Geo-replication
 - o Volume Split-Brain
 - o Cluster Quorum (A cluster quorum service would be present only when there are volumes in the cluster.)

Rescheduling Cluster Auto config using Nagios Server GUI

If new nodes or services are added or removed, or if snapshot restore is performed on Red Hat Gluster Storage node, reschedule the **Cluster Auto config** service using Nagios Server GUI or execute the **configure-gluster-nagios** command. To synchronize the configurations using Nagios Server GUI, perform the steps given below:

1. Login to the Nagios Server GUI using the following URL in your browser with nagiosadmin user name and password.

<https://NagiosServer-HostName-or-IPaddress/nagios>

2. Click **Services** in left pane of Nagios server GUI and click **Cluster Auto Config**.

The screenshot shows the Nagios web interface. On the left is a navigation menu with 'Services' selected. The main content area displays 'Service Status Details For All Host Groups'. At the top, there are summary statistics for Host Status Totals (Up: 2, Down: 1, Unreachable: 0, Pending: 0) and Service Status Totals (Ok: 29, Warning: 2, Unknown: 4, Critical: 11, Pending: 0). Below these is a table of service details. The 'Cluster Auto Config' service is highlighted in red, indicating a critical status. The table columns include Host, Service, Status, Last Check, Duration, Attempt, and Status Information.

Host	Service	Status	Last Check	Duration	Attempt	Status Information
cluster1	Cluster - Quorum Status	UNKNOWN	07-25-2015 12:18:58	1d 18h 36m 44s	1/1	Server quorum not turned on for any volume
	Cluster Auto Config	OK	07-24-2015 13:18:58	0d 23h 58m 35s	1/3	Cluster configurations are in sync
	Cluster Utilization	UNKNOWN	07-25-2015 13:07:46	0d 16h 17m 44s	1/3	UNKNOWN - Volume utilization data could not be read
	Volume Geo-Replication - vol1	OK	07-25-2015 13:10:58	0d 16h 4m 32s	1/3	Session status: No active sessions found
	Volume Split-brain status - gluster_shared_storage	WARNING	07-25-2015 13:06:58	0d 16h 22m 32s	3/3	Volume split brain status could not be determined.
	Volume Split-brain status - vol1	WARNING	07-25-2015 13:08:58	0d 16h 16m 32s	3/3	Volume split brain status could not be determined.
	Volume Status - gluster_shared_storage	CRITICAL	07-25-2015 13:08:58	0d 16h 16m 32s	2/3	CRITICAL: Volume not found
	Volume Status - vol1	CRITICAL	07-25-2015 13:12:58	0d 16h 26m 32s	3/3	CRITICAL: Volume - DISTRIBUTE type
	Volume Status - vol_dis_rep	CRITICAL	07-25-2015 13:12:46	0d 16h 26m 44s	3/3	CRITICAL: Volume not found
	Volume Utilization - gluster_shared_storage	CRITICAL	07-25-2015 13:14:58	0d 16h 24m 32s	3/3	CRITICAL: Volume not found
rhs-client24.lab.eng.bjr.redhat.com	Volume Utilization - vol1	CRITICAL	07-25-2015 13:06:58	0d 16h 22m 32s	3/3	CRITICAL: Failed to get the Volume Utilization Data
	Volume Utilization - vol_dis_rep	CRITICAL	07-25-2015 13:13:46	0d 16h 25m 44s	3/3	CRITICAL: Volume not found
	Brick - /rhgs/brick1/brick1	CRITICAL	07-25-2015 13:14:20	0d 16h 5m 10s	3/3	Brick - /rhgs/brick1/brick1 not found
	Brick - /rhgs/brick2/brick2	CRITICAL	07-25-2015 13:10:21	0d 16h 27m 9s	3/3	Brick - /rhgs/brick2/brick2 not found
	Brick - /var/lib/glusterd/sss_brick	CRITICAL	07-25-2015 13:10:22	0d 16h 27m 8s	3/3	Brick - /var/lib/glusterd/sss_brick not found
	Brick Utilization - /rhgs/brick1/brick1	OK	07-25-2015 13:05:46	1d 19h 19m 44s	1/3	OK:
	Brick Utilization - /rhgs/brick2/brick2	OK	07-25-2015 13:06:58	0d 23h 36m 32s	1/3	OK:
	Brick Utilization - /var/lib/glusterd/sss_brick	OK	07-25-2015 13:10:58	0d 23h 34m 32s	1/3	OK:
	CTDB	UNKNOWN	07-24-2015 17:58:46	1d 19h 26m 26s	1/3	CTDB not configured
	Cpu Utilization	OK	07-25-2015 13:14:46	1d 19h 13m 44s	3/3	CPU Status OK: Total CPU:0.34% Idle CPU:99.66%
Disk Utilization	OK	07-25-2015 13:14:46	1d 19h 10m 44s	1/3	OK:	
Gluster Management	OK	07-24-2015 18:07:46	1d 19h 26m 29s	1/3	Process glusterd is running	

Figure 13.10. Nagios Services

3. In **Service Commands**, click **Re-schedule the next check of this service**. The **Command Options** window is displayed.

The screenshot shows the 'Service Commands' dialog box. It contains a list of actions with corresponding icons. The 'Re-schedule the next check of this service' option, which has a clock icon, is highlighted with a red rectangular box. Other options include 'Disable active checks of this service', 'Submit passive check result for this service', 'Stop accepting passive checks for this service', 'Stop obsessing over this service', 'Disable notifications for this service', 'Send custom service notification', 'Schedule downtime for this service', 'Disable event handler for this service', and 'Disable flap detection for this service'.

Figure 13.11. Service Commands

4. In **Command Options** window, click **Commit**.

You are requesting to schedule a service check

Command Options		Command Description
Host Name:	<input type="text" value="test-cluster"/>	This command is used to schedule the next check of a particular service. Nagios will re-queue the service to be checked at the time you specify. If you select the <i>force check</i> option, Nagios will force a check of the service regardless of both what time the scheduled check occurs and whether or not checks are enabled for the service.
Service:	<input type="text" value="Cluster Auto Config"/>	
Check Time:	<input type="text" value="06-26-2014 14:36:13"/>	
Force Check:	<input checked="" type="checkbox"/>	
<input type="button" value="Commit"/> <input type="button" value="Reset"/>		

Please enter all required information before committing the command.
 Required fields are marked in red.
 Failure to supply all required values will result in an error.

Figure 13.12. Command Options

Enabling and Disabling Notifications using Nagios GUI

You can enable or disable Host and Service notifications through Nagios GUI.

- To enable and disable Host Notifications:
 1. Login to the Nagios Server GUI using the following URL in your browser with nagiosadmin user name and password.

`https://NagiosServer-HostName-or-IPaddress/nagios`
 2. Click **Hosts** in left pane of Nagios server GUI and select the host.
 3. Click **Enable notifications for this host** or **Disable notifications for this host** in Host Commands section.
 4. Click **Commit** to enable or disable notification for the selected host.
- To enable and disable Service Notification:
 1. Login to the Nagios Server GUI.
 2. Click **Services** in left pane of Nagios server GUI and select the service to enable or disable.
 3. Click **Enable notifications for this service** or **Disable notifications for this service** from the Service Commands section.
 4. Click **Commit** to enable or disable the selected service notification.
- To enable and disable all Service Notifications for a host:
 1. Login to the Nagios Server GUI.
 2. Click **Hosts** in left pane of Nagios server GUI and select the host to enable or disable all services notifications.
 3. Click **Enable notifications for all services on this host** or **Disable notifications for all services on this host** from the Service Commands section.

4. Click **Commit** to enable or disable all service notifications for the selected host.
- To enable or disable all Notifications:
 1. Login to the Nagios Server GUI.
 2. Click **Process Info** under **Systems** section from left pane of Nagios server GUI.
 3. Click **Enable notifications** or **Disable notifications** in Process Commands section.
 4. Click **Commit**.

Enabling and Disabling Service Monitoring using Nagios GUI

You can enable a service to monitor or disable a service you have been monitoring using the Nagios GUI.

- To enable Service Monitoring:
 1. Login to the Nagios Server GUI using the following URL in your browser with nagiosadmin user name and password.

```
https://NagiosServer-HostName-or-IPaddress/nagios
```
 2. Click **Services** in left pane of Nagios server GUI and select the service to enable monitoring.
 3. Click **Enable active checks of this service** from the Service Commands and click **Commit**.
 4. Click **Start accepting passive checks for this service** from the Service Commands and click **Commit**.

Monitoring is enabled for the selected service.

- To disable Service Monitoring:
 1. Login to the Nagios Server GUI using the following URL in your browser with nagiosadmin user name and password.

```
https://NagiosServer-HostName-or-IPaddress/nagios
```
 2. Click **Services** in left pane of Nagios server GUI and select the service to disable monitoring.
 3. Click **Disable active checks of this service** from the Service Commands and click **Commit**.
 4. Click **Stop accepting passive checks for this service** from the Service Commands and click **Commit**.

Monitoring is disabled for the selected service.

Monitoring Services Status and Messages

**NOTE**

Nagios sends email and SNMP notifications, once a service status changes. Refer *Configuring Nagios Server to Send Mail Notifications* section of *Red Hat Gluster Storage 3.4 Console Administration Guide* to configure email notification and *Configuring Simple Network Management Protocol (SNMP) Notification* section of *Red Hat Gluster Storage 3.4 Administration Guide* to configure SNMP notification.

Table 13.1.

Service Name	Status	Message	Description
SMB	OK	OK: No gluster volume uses smb	When no volumes are exported through smb.
	OK	Process smb is running	When SMB service is running and when volumes are exported using SMB.
	CRITICAL	CRITICAL: Process smb is not running	When SMB service is down and one or more volumes are exported through SMB.
CTDB	UNKNOWN	CTDB not configured	When CTDB service is not running, and smb or nfs service is running.
	CRITICAL	Node status: BANNED/STOPPED	When CTDB service is running but Node status is <i>BANNED/STOPPED</i> .
	WARNING	Node status: UNHEALTHY/DISABLED/PARTIALLY_ONLINE	When CTDB service is running but Node status is <i>UNHEALTHY/DISABLED/PARTIALLY_ONLINE</i> .
	OK	Node status: OK	When CTDB service is running and healthy.
Gluster Management	OK	Process glusterd is running	When glusterd is running as unique.
	WARNING	PROCS WARNING: 3 processes	When there are more than one glusterd is running.

Service Name	Status	Message	Description
	CRITICAL	CRITICAL: Process glusterd is not running	When there is no glusterd process running.
	UNKNOWN	NRPE: Unable to read output	When unable to communicate or read output
Gluster NFS	OK	OK: No gluster volume uses nfs	When no volumes are configured to be exported through NFS.
	OK	Process glusterfs-nfs is running	When glusterfs-nfs process is running.
	CRITICAL	CRITICAL: Process glusterfs-nfs is not running	When glusterfs-nfs process is down and there are volumes which requires NFS export.
Split-brain	OK	No Split brain entries found.	When there are files present with out any split brain issues.
	WARNING	Volume split brain status could not be determined	
	CRITICAL	CRITICAL: <i>No.of files</i> in split brain state found.	When there are some files in split brain state.
Auto-Config	OK	Cluster configurations are in sync	When auto-config has not detected any change in Gluster configuration. This shows that Nagios configuration is already in synchronization with the Gluster configuration and auto-config service has not made any change in Nagios configuration.

Service Name	Status	Message	Description
	OK	Cluster configurations synchronized successfully from host <i>host-address</i>	When auto-config has detected change in the Gluster configuration and has successfully updated the Nagios configuration to reflect the change Gluster configuration.
	CRITICAL	Can't remove all hosts except sync host in 'auto' mode. Run auto discovery manually.	When the host used for auto-config itself is removed from the Gluster peer list. Auto-config will detect this as all host except the synchronized host is removed from the cluster. This will not change the Nagios configuration and the user need to manually run the auto-config.
QUOTA	OK	OK: Quota not enabled	When quota is not enabled in any volumes.
	OK	Process quotad is running	When glusterfs-quota service is running.
	CRITICAL	CRITICAL: Process quotad is not running	When glusterfs-quota service is down and quota is enabled for one or more volumes.
CPU Utilization	OK	CPU Status OK: Total CPU:4.6% Idle CPU:95.40%	When CPU usage is less than 80%.
	WARNING	CPU Status WARNING: Total CPU:82.40% Idle CPU:17.60%	When CPU usage is more than 80%.
	CRITICAL	CPU Status CRITICAL: Total CPU:97.40% Idle CPU:2.6%	When CPU usage is more than 90%.

Service Name	Status	Message	Description
Memory Utilization	OK	OK- 65.49% used(1.28GB out of 1.96GB)	When used memory is below warning threshold. (Default warning threshold is 80%)
	WARNING	WARNING- 85% used(1.78GB out of 2.10GB)	When used memory is below critical threshold (Default critical threshold is 90%) and greater than or equal to warning threshold (Default warning threshold is 80%).
	CRITICAL	CRITICAL- 92% used(1.93GB out of 2.10GB)	When used memory is greater than or equal to critical threshold (Default critical threshold is 90%)
Brick Utilization	OK	OK	When used space of any of the four parameters, space detail, inode detail, thin pool, and thin pool-metadata utilizations, are below threshold of 80%.
	WARNING	WARNING:mount point /brick/brk1 Space used (0.857 / 1.000) GB	If any of the four parameters, space detail, inode detail, thin pool utilization, and thinpool-metadata utilization, crosses warning threshold of 80% (Default is 80%).
	CRITICAL	CRITICAL : mount point /brick/brk1 (inode used 9980/1000)	If any of the four parameters, space detail, inode detail, thin pool utilization, and thinpool-metadata utilizations, crosses critical threshold 90% (Default is 90%).

Service Name	Status	Message	Description
Disk Utilization	OK	OK	When used space of any of the four parameters, space detail, inode detail, thin pool utilization, and thinpool-metadata utilizations, are below threshold of 80%.
	WARNING	WARNING:mount point /boot Space used (0.857 / 1.000) GB	When used space of any of the four parameters, space detail, inode detail, thin pool utilization, and thinpool-metadata utilizations, are above warning threshold of 80%.
	CRITICAL	CRITICAL : mount point /home (inode used 9980/1000)	If any of the four parameters, space detail, inode detail, thin pool utilization, and thinpool-metadata utilizations, crosses critical threshold 90% (Default is 90%).
Network Utilization	OK	OK: tun0:UP,wlp3s0:UP,virbr0:UP	When all the interfaces are UP.
	WARNING	WARNING: tun0:UP,wlp3s0:UP,virbr0:DOWN	When any of the interfaces is down.
	UNKNOWN	UNKNOWN	When network utilization/status is unknown.
Swap Utilization	OK	OK- 0.00% used(0.00GB out of 1.00GB)	When used memory is below warning threshold (Default warning threshold is 80%).

Service Name	Status	Message	Description
	WARNING	WARNING- 83% used(1.24GB out of 1.50GB)	When used memory is below critical threshold (Default critical threshold is 90%) and greater than or equal to warning threshold (Default warning threshold is 80%).
	CRITICAL	CRITICAL- 83% used(1.42GB out of 1.50GB)	When used memory is greater than or equal to critical threshold (Default critical threshold is 90%).
Cluster- Quorum	PENDING		When cluster.quorum-type is not set to <i>server</i> ; or when there are no problems in the cluster identified.
	OK	Quorum regained for volume	When quorum is regained for volume.
	CRITICAL	Quorum lost for volume	When quorum is lost for volume.
Volume Geo-replication	OK	"Session Status: <i>slave_vol1</i> -OK <i>slave_voln</i> -OK.	When all sessions are active.
		Session status :No active sessions found	When Geo-replication sessions are deleted.
	CRITICAL	Session Status: <i>slave_vol1</i> -FAULTY <i>slave_vol2</i> -OK	If one or more nodes are Faulty and there's no replica pair that's active.

Service Name	Status	Message	Description
	WARNING	Session Status: <i>slave_vol1</i> - NOT_STARTED <i>slave_vol2</i> -STOPPED <i>slave_vol3</i> - PARTIAL_FAULTY	<ul style="list-style-type: none"> • Partial faulty state occurs with replicated and distributed replicate volume when one node is faulty, but the replica pair is active. • <i>STOPPED</i> state occurs when Geo-replication sessions are stopped. • <i>NOT_STARTED</i> state occurs when there are multiple Geo-replication sessions and one of them is stopped.
	WARNING	Geo replication status could not be determined.	When there's an error in getting Geo replication status. This error occurs when volfile is locked as another transaction is in progress.
	UNKNOWN	Geo replication status could not be determined.	When glusterd is down.
Volume Quota	OK	QUOTA: not enabled or configured	When quota is not set
	OK	QUOTA:OK	When quota is set and usage is below quota limits.
	WARNING	QUOTA:Soft limit exceeded on <i>path of directory</i>	When quota exceeds soft limit.
	CRITICAL	QUOTA:hard limit reached on <i>path of directory</i>	When quota reaches hard limit.

Service Name	Status	Message	Description
	UNKNOWN	QUOTA: Quota status could not be determined as command execution failed	When there's an error in getting Quota status. This occurs when <ul style="list-style-type: none"> • Volume is stopped or glusterd service is down. • volfile is locked as another transaction in progress.
Volume Status	OK	Volume : <i>volume type</i> - All bricks are Up	When all volumes are up.
	WARNING	Volume : <i>volume type</i> Brick(s) - <i>list of bricks</i> is are down, but replica pair(s) are up	When bricks in the volume are down but replica pairs are up.
	UNKNOWN	Command execution failed <i>Failure message</i>	When command execution fails.
	CRITICAL	Volume not found.	When volumes are not found.
	CRITICAL	Volume: <i>volume-type</i> is stopped.	When volumes are stopped.
	CRITICAL	Volume : <i>volume type</i> - All bricks are down.	When all bricks are down.
	CRITICAL	Volume : <i>volume type</i> Bricks - <i>brick list</i> are down, along with one or more replica pairs	When bricks are down along with one or more replica pairs.
Volume Self-Heal	OK		When volume is not a replicated volume, there is no self-heal to be done.
	OK	No unsynched entries present	When there are no unsynched entries in a replicated volume.

Service Name	Status	Message	Description
	WARNING	Unsynced entries present : There are unsynced entries present.	If self-heal process is turned on, these entries may be auto healed. If not, self-heal will need to be run manually. If unsynchronized entries persist over time, this could indicate a split brain scenario.
	WARNING	Self heal status could not be determined as the volume was deleted	When self-heal status can not be determined as the volume is deleted.
	UNKNOWN		When there's an error in getting self heal status. This error occurs when: <ul style="list-style-type: none"> • Volume is stopped or glusterd service is down. • volfile is locked as another transaction in progress.
Cluster Utilization	OK	OK : 28.0% used (1.68GB out of 6.0GB)	When used % is below the warning threshold (Default warning threshold is 80%).
	WARNING	WARNING: 82.0% used (4.92GB out of 6.0GB)	Used% is above the warning limit. (Default warning threshold is 80%)
	CRITICAL	CRITICAL : 92.0% used (5.52GB out of 6.0GB)	Used% is above the warning limit. (Default critical threshold is 90%)
	UNKNOWN	Volume utilization data could not be read	When volume services are present, but the volume utilization data is not available as it's either not populated yet or there is error in fetching volume utilization data.

Service Name	Status	Message	Description
Volume Utilization	OK	OK: Utilization: 40 %	When used % is below the warning threshold (Default warning threshold is 80%).
	WARNING	WARNING - used 84% of available 200 GB	When used % is above the warning threshold (Default warning threshold is 80%).
	CRITICAL	CRITICAL - used 96% of available 200 GB	When used % is above the critical threshold (Default critical threshold is 90%).
	UNKNOWN	UNKNOWN - Volume utilization data could not be read	When all the bricks in the volume are killed or if glusterd is stopped in all the nodes in a cluster.

13.5. MONITORING HOST AND CLUSTER UTILIZATION

13.5.1. Monitoring Host and Cluster Utilization

You can monitor utilization and set alerts and notifications for the utilization changes.

You can monitor Host and Cluster utilization using Nagios plug-in and validate the status of clusters and hosts from the utilization graph through Red Hat Gluster Storage Console.



NOTE

By default, you can view the Utilization report of the last 24 hours.

Procedure 13.1. To Monitor Cluster Utilization

1. Click **System** and select **Clusters** in the Tree pane.
2. Click **Trends** tab.

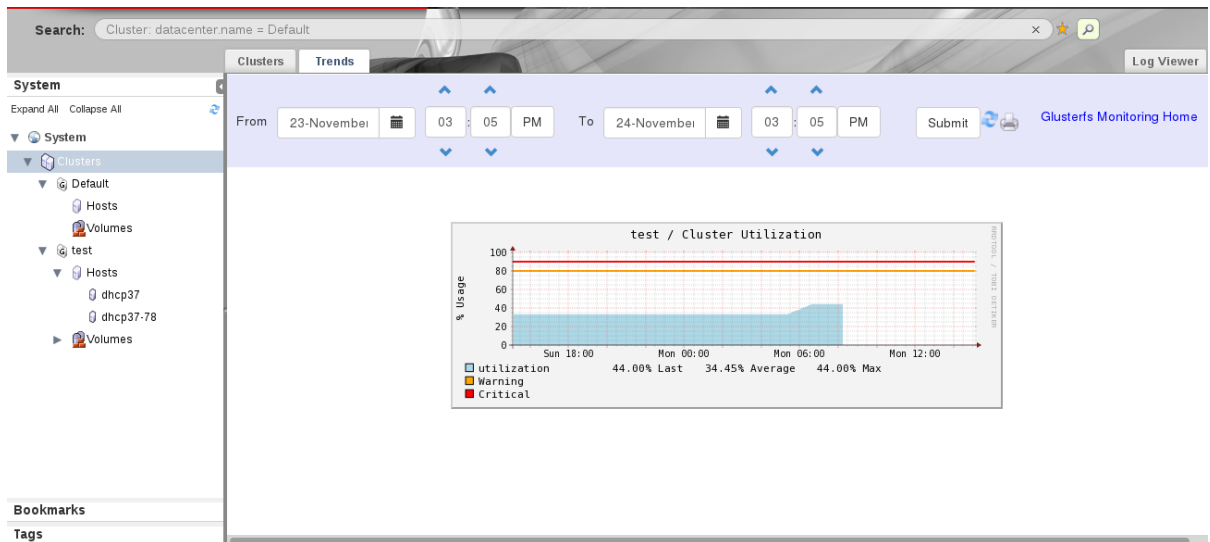


Figure 13.13. Trends

3. Select the date and time duration to view the cluster utilization report.
4. Click **Submit**. The Cluster Utilization graph of all clusters for the selected period is displayed.

You can refresh the status by clicking the refresh button and also print the report or save as a pdf file by clicking the print button. Click **Glusterfs Monitoring Home** to view the Nagios Home page.

Procedure 13.2. To Monitor Utilization for Hosts

1. Click **System** and select **Clusters** in the Tree pane.
2. Click **Hosts** in the tree pane and click **Trends** tab to view the CPU Utilization for all the hosts.

To view *CPU Utilization*, *Network Interface Utilization*, *Disk Utilization*, *Memory Utilization* and *Swap Utilization* for each host, select the Host name from the tree pane and click **Trends** tab.

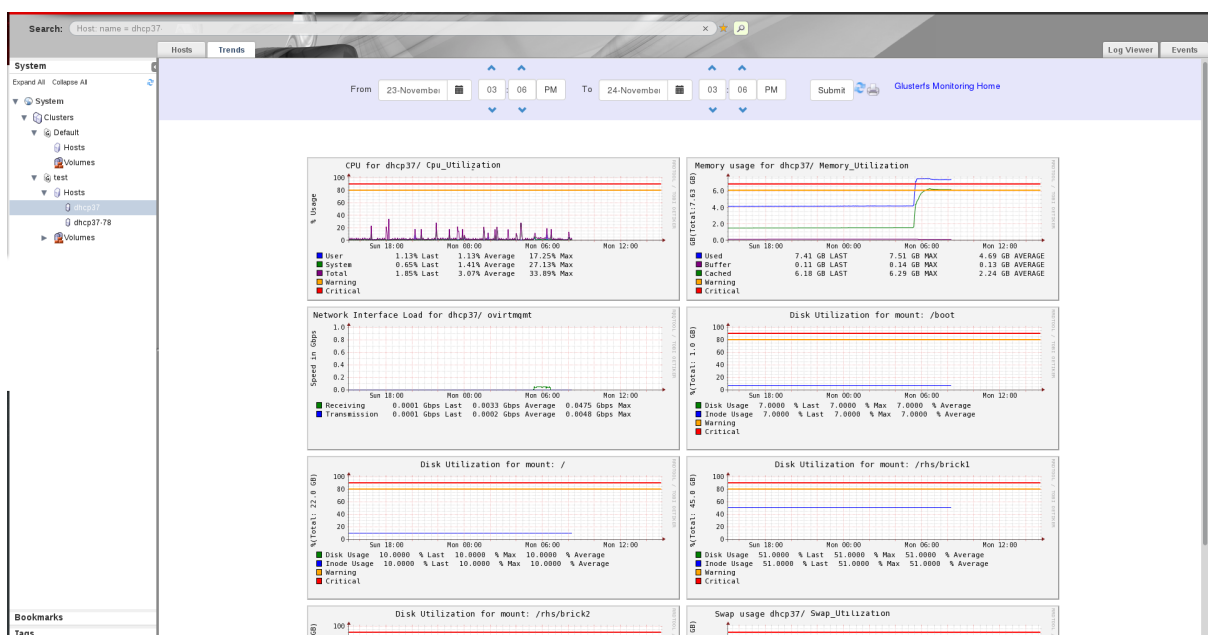


Figure 13.14. Utilization for selected Host

3. Select the date and time to view the Host Utilization report.
4. Click **Submit**. The CPU Utilization graph for all the Hosts for the selected period is displayed.

You can refresh the status by clicking the refresh button and also print the report or save as a pdf file by clicking the print button. To view the Nagios Home page, click **Glusterfs Monitoring Home**.

Procedure 13.3. To monitor Volume and Brick Utilization

1. Open the **Volumes** view in the tree pane and select **Volumes**.
2. Click **Trends** tab.
3. Select the date and time duration to view the volume and brick utilization report.
4. Click **Submit**. The Volume Utilization graph and Brick Utilization graph for the selected period is displayed.

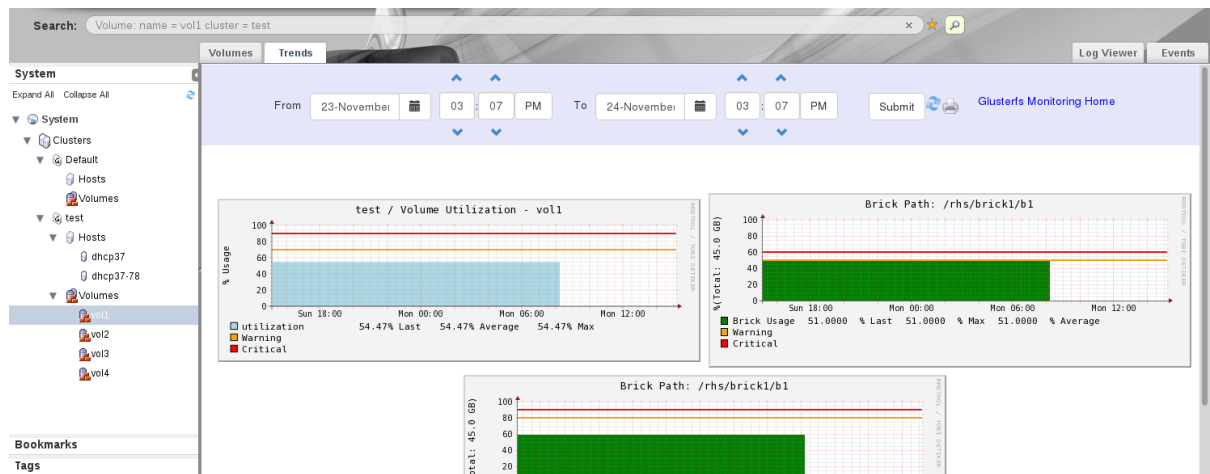



Figure 13.15. Volume and Brick Utilization

You can refresh the status by clicking the refresh button and also print the report or save as a pdf file by clicking the print button. To view the Nagios Home page, click **Glusterfs Monitoring Home**.

13.5.2. Enabling and Disabling Monitoring

You can enable and disable monitoring using the command line interface after setting up Red Hat Gluster Storage Console. The **Trends** tab is displayed with host and cluster utilization details when monitoring is enabled on the server.



IMPORTANT

You must refresh the browser after enabling or disabling monitoring to view the changes.

- To enable monitoring, run the following command in the Red Hat Gluster Storage Console Server :

```
# rhsc-monitoring enable
Setting the monitoring flag...
```

```
Starting nagios: done.
Starting nsca: [ OK ]
INFO: Move the nodes of existing cluster (with compatibility version
>= 3.4) to maintenance and re-install them.
```

The **Trends** tab is displayed in the Red Hat Gluster Storage Console Administrator portal with the host and cluster utilization details.

- To disable monitoring, run the following command in the Red Hat Gluster Storage Console Server:

```
# rhsc-monitoring disable
Setting the monitoring flag...
Stopping nagios: .done.
Shutting down nsca: [ OK ]
```

The **Trends** tab is not displayed in the Red Hat Gluster Storage Console Administrator portal and the user cannot view host and cluster utilization details. Receiving email and SNMP notifications are disabled. Disabling monitoring also stops Nagios and NSCA services.

Disabling monitoring does not stop the **glusterpmd** service. Run the following commands on all the Red Hat Gluster Storage nodes to stop **glusterpmd** service and to remove **chkconfig** for **glusterpmd** service:

```
# service glusterpmd stop
# chkconfig glusterpmd off
```

13.6. TROUBLESHOOTING NAGIOS

13.6.1. Troubleshooting NSCA and NRPE Configuration Issues

The possible errors while configuring Nagios Service Check Acceptor (NSCA) and Nagios Remote Plug-in Executor (NRPE) and the troubleshooting steps are listed in this section.

Troubleshooting NSCA Configuration Issues

- **Check Firewall and Port Settings on Nagios Server**

If port 5667 is not opened on the server host's firewall, a timeout error is displayed. Ensure that port 5667 is opened.

On Red Hat Gluster Storage based on Red Hat Enterprise Linux 6

1. Log in as root and run the following command on the Red Hat Gluster Storage node to get the list of current iptables rules:

```
# iptables -L
```

2. The output is displayed as shown below:

```
ACCEPT      tcp  --  anywhere          anywhere          tcp
dpt:5667
```

On Red Hat Gluster Storage based on Red Hat Enterprise Linux 7:

1. Run the following command on the Red Hat Gluster Storage node as root to get a listing of the current firewall rules:

```
# firewall-cmd --list-all-zones
```

2. If the port is open, **5667/tcp** is listed beside **ports:** under one or more zones in your output.
- If the port is not open, add a firewall rule for the port:

On Red Hat Gluster Storage based on Red Hat Enterprise Linux 6

1. If the port is not open, add an iptables rule by adding the following line in **/etc/sysconfig/iptables** file:

```
-A INPUT -m state --state NEW -m tcp -p tcp --dport 5667 -j  
ACCEPT
```

2. Restart the iptables service using the following command:

```
# service iptables restart
```

3. Restart the NSCA service using the following command:

```
# service nsca restart
```

On Red Hat Gluster Storage based on Red Hat Enterprise Linux 7:

1. Run the following commands to open the port:

```
# firewall-cmd --zone=public --add-port=5667/tcp  
# firewall-cmd --zone=public --add-port=5667/tcp --permanent
```

- **Check the Configuration File on Red Hat Gluster Storage Node**

Messages cannot be sent to the NSCA server, if Nagios server IP or FQDN, cluster name and hostname (as configured in Nagios server) are not configured correctly.

Open the Nagios server configuration file **/etc/nagios/nagios_server.conf** and verify if the correct configurations are set as shown below:

```
# NAGIOS SERVER  
# The nagios server IP address or FQDN to which the NSCA command  
# needs to be sent  
[NAGIOS-SERVER]  
nagios_server=NagiosServerIPAddress  
  
# CLUSTER NAME  
# The host name of the logical cluster configured in Nagios under  
which
```



```
# the gluster volume services reside
[NAGIOS-DEFINTIONS]
cluster_name=cluster_auto

# LOCAL HOST NAME
# Host name given in the nagios server
[HOST-NAME]
hostname_in_nagios=NagiosServerHostName
```

If Host name is updated, restart the NSCA service using the following command:

```
# service nsca restart
```

Troubleshooting NRPE Configuration Issues

- **CHECK_NRPE: Error - Could Not Complete SSL Handshake**

This error occurs if the IP address of the Nagios server is not defined in the `nrpe.cfg` file of the Red Hat Gluster Storage node. To fix this issue, follow the steps given below:

1. Add the Nagios server IP address in `/etc/nagios/nrpe.cfg` file in the `allowed_hosts` line as shown below:

```
allowed_hosts=127.0.0.1, NagiosServerIP
```

The `allowed_hosts` is the list of IP addresses which can execute NRPE commands.

2. Save the `nrpe.cfg` file and restart the NRPE service using the following command:

```
# service nrpe restart
```

- **CHECK_NRPE: Socket Timeout After n Seconds**

To resolve this issue perform the steps given below:

On Nagios Server:

The default timeout value for the NRPE calls is 10 seconds and if the server does not respond within 10 seconds, Nagios GUI displays an error that the NRPE call has timed out in 10 seconds. To fix this issue, change the timeout value for NRPE calls by modifying the command definition configuration files.

1. Changing the NRPE timeout for services which directly invoke `check_nrpe`.

For the services which directly invoke `check_nrpe` (`check_disk_and_inode`, `check_cpu_multicore`, and `check_memory`), modify the command definition configuration file `/etc/nagios/gluster/gluster-commands.cfg` by adding `-t Time in Seconds` as shown below:

```
define command {
    command_name check_disk_and_inode
    command_line $USER1$/check_nrpe -H $HOSTADDRESS$ -c
check_disk_and_inode -t TimeInSeconds
}
```

2. Changing the NRPE timeout for the services in **nagios-server-addons** package which invoke NRPE call through code.

The services which invoke

/usr/lib64/nagios/plugins/gluster/check_vol_server.py

(`check_vol_utilization`, `check_vol_status`, `check_vol_quota_status`, `check_vol_heal_status`, and `check_vol_georep_status`) make NRPE call to the Red Hat Gluster Storage nodes for the details through code. To change the timeout for the NRPE calls, modify the command definition configuration file **/etc/nagios/gluster/gluster-commands.cfg** by adding `-t No of seconds` as shown below:

```
define command {
    command_name check_vol_utilization
    command_line $USER1$/gluster/check_vol_server.py $ARG1$
    $ARG2$ -w $ARG3$ -c $ARG4$ -o utilization -t TimeInSeconds
}
```

The auto configuration service **gluster_auto_discovery** makes NRPE calls for the configuration details from the Red Hat Gluster Storage nodes. To change the NRPE timeout value for the auto configuration service, modify the command definition configuration file **/etc/nagios/gluster/gluster-commands.cfg** by adding `-t TimeInSeconds` as shown below:

```
define command{
    command_name    gluster_auto_discovery
    command_line    sudo $USER1$/gluster/configure-gluster-
nagios.py -H $ARG1$ -c $HOSTNAME$ -m auto -n $ARG2$ -t
TimeInSeconds
}
```

3. Restart Nagios service using the following command:

```
# service nagios restart
```

On Red Hat Gluster Storage node:

1. Add the Nagios server IP address as described in *CHECK_NRPE: Error - Could Not Complete SSL Handshake* section in *Troubleshooting NRPE Configuration Issues* section.
2. Edit the **nrpe.cfg** file using the following command:

```
# vi /etc/nagios/nrpe.cfg
```

3. Search for the **command_timeout** and **connection_timeout** settings and change the value. The **command_timeout** value must be greater than or equal to the timeout value set in Nagios server.

The timeout on checks can be set as `connection_timeout=300` and the `command_timeout=60` seconds.

4. Restart the NRPE service using the following command:

```
# service nrpe restart
```

- **Check the NRPE Service Status**

This error occurs if the NRPE service is not running. To resolve this issue perform the steps given below:

1. Verify the status of NRPE service by logging into the Red Hat Gluster Storage node as root and running the following command:

```
# service nrpe status
```

2. If NRPE is not running, start the service using the following command:

```
# service nrpe start
```

- **Check Firewall and Port Settings**

This error is associated with firewalls and ports. The timeout error is displayed if the NRPE traffic is not traversing a firewall, or if port 5666 is not open on the Red Hat Gluster Storage node.

Ensure that port 5666 is open on the Red Hat Gluster Storage node.

1. Run **check_nrpe** command from the Nagios server to verify if the port is open and if NRPE is running on the Red Hat Gluster Storage Node .
2. Log into the Nagios server as root and run the following command:

```
# /usr/lib64/nagios/plugins/check_nrpe -H RedHatStorageNodeIP
```

3. The output is displayed as given below:

```
NRPE v2.14
```

If not, ensure the that port 5666 is opened on the Red Hat Gluster Storage node.

On Red Hat Gluster Storage based on Red Hat Enterprise Linux 6:

1. Run the following command on the Red Hat Gluster Storage node as root to get a listing of the current iptables rules:

```
# iptables -L
```

2. If the port is open, the following appears in your output.

```
ACCEPT      tcp  --  anywhere                anywhere                tcp
dpt:5666
```

On Red Hat Gluster Storage based on Red Hat Enterprise Linux 7:

1. Run the following command on the Red Hat Gluster Storage node as root to get a listing of the current firewall rules:

```
# firewall-cmd --list-all-zones
```

-
- 2. If the port is open, **5666/tcp** is listed beside **ports:** under one or more zones in your output.
- If the port is not open, add an iptables rule for the port.

On Red Hat Gluster Storage based on Red Hat Enterprise Linux 6:

1. To add iptables rule, edit the **iptables** file as shown below:

```
# vi /etc/sysconfig/iptables
```

2. Add the following line in the file:

```
-A INPUT -m state --state NEW -m tcp -p tcp --dport 5666 -j  
ACCEPT
```

3. Restart the iptables service using the following command:

```
# service iptables restart
```

4. Save the file and restart the NRPE service:

```
# service nrpe restart
```

On Red Hat Gluster Storage based on Red Hat Enterprise Linux 7:

1. Run the following commands to open the port:

```
# firewall-cmd --zone=public --add-port=5666/tcp  
# firewall-cmd --zone=public --add-port=5666/tcp --permanent
```

● Checking Port 5666 From the Nagios Server with Telnet

Use telnet to verify the Red Hat Gluster Storage node's ports. To verify the ports of the Red Hat Gluster Storage node, perform the steps given below:

1. Log in as root on Nagios server.
2. Test the connection on port 5666 from the Nagios server to the Red Hat Gluster Storage node using the following command:

```
# telnet RedHatStorageNodeIP 5666
```

3. The output displayed is similar to:

```
telnet 10.70.36.49 5666  
Trying 10.70.36.49...  
Connected to 10.70.36.49.  
Escape character is '^]'.  
^C
```

● Connection Refused By Host

This error is due to port/firewall issues or incorrectly configured *allowed_hosts* directives. See the sections *CHECK_NRPE: Error - Could Not Complete SSL Handshake* and *CHECK_NRPE: Socket Timeout After n Seconds* for troubleshooting steps.

13.6.2. Troubleshooting General Issues

This section describes the troubleshooting procedures for general issues related to Nagios.

Graphs are not displayed in Trends tab

Ensure that the host name given in Name field of Add Host window matches the host name given while configuring Nagios. The host name of the node is used while configuring Nagios server using auto-discovery.

PART IV. MANAGING ADVANCED FUNCTIONALITY

CHAPTER 14. MANAGING MULTILEVEL ADMINISTRATION

Red Hat Gluster Storage Console supports multilevel administration. That is, users can be assigned a variety of permissions for specific objects using a number of default roles. This section describes how to set up user roles that control levels of permissions for different objects and actions in your storage environment. Customized roles can also be created and assigned to users.

Red Hat Gluster Storage Console relies on directory services for user authentication. The providers of directory services currently supported for use with the Red Hat Gluster Storage Console are Identity (IdM), Active Directory, and Red Hat Directory Server (RHDS).



NOTE

Users are not created in Red Hat Gluster Storage, but in the Directory Services domain. Red Hat Gluster Storage Console can be configured to use multiple Directory Services domains. See the *Red Hat Gluster Storage Console Installation Guide* for more information.

14.1. CONFIGURING ROLES

Roles are predefined sets of privileges that can be configured from Red Hat Gluster Storage Console, providing access and management permissions to different levels of resources in the cluster. Permissions enable users to perform actions on objects.

With multilevel administration, any permissions that apply to a container object also apply to all individual objects within that container. For example, when a server administrator role is assigned to a user on a specific server, the user gains permissions to perform any of the available operations, but only on the assigned server. However, if the administrator role is assigned to a user on a cluster, the user gains permissions to perform operations on all servers within the cluster.

14.1.1. Roles

There is one type of role in Red Hat Gluster Storage Console, which is the **administrator** role. This role allows access to the Administration Portal for managing server resources. For example, if a user has an **administrator** role on a cluster, they can manage all servers in the cluster using the Administration Portal.

The default roles cannot be removed from the Red Hat Gluster Storage, and their privileges cannot be modified. However, you can clone them and then customize the new roles as required.

14.1.2. Creating Custom Roles

In addition to the default roles, you can set up custom roles that permit actions on objects, such as servers and clusters, and assign privileges to specific entities. Use roles to create a granular model of permissions to suit the needs of the enterprise or a group or set of users. Use the **Configure** option to work with roles. You can create a **New** role, or **Edit**, **Clone** or **Remove** an existing role. In each case, the appropriate dialog box displays.

Once the role is set up, you can assign the role to users as required.

Procedure 14.1. Creating a New Role

1. On the header bar of the Red Hat Gluster Storage Console menu, click **Configure**. The **Configure** dialog box displays. The dialog box includes a list of Administrator roles, and any custom roles.
2. Click **New**. The **New Role** dialog box displays.
3. Enter the **Name** and **Description** of the new role. This name will display in the list of roles.
4. Select **Admin** as the **Account Type**. If **Admin** is selected, this role displays with the administrator icon in the list.
5. Use the **Expand All** or **Collapse All** buttons to view more or fewer of the permissions for the listed objects in the **Check Boxes to Allow Action** list. You can also expand or collapse the options for each object.
6. For each of the objects, select or deselect the actions you wish to permit/deny for the role you are setting up.
7. Click **OK** to apply the changes you have made. The new role displays on the list of roles.

14.1.3. Editing Roles

While you cannot make changes to the default roles, you may need to change the permissions, names or descriptions of custom roles. To edit custom roles, use the **Edit** button on the **Configure** dialog box.

Procedure 14.2. Editing a Role

1. On the header bar of the Red Hat Gluster Storage Console menu, click **Configure**. The **Configure** dialog box displays. The dialog box below shows the list of administrator roles.
2. Click **Edit**. The **Edit Role** dialog box displays.

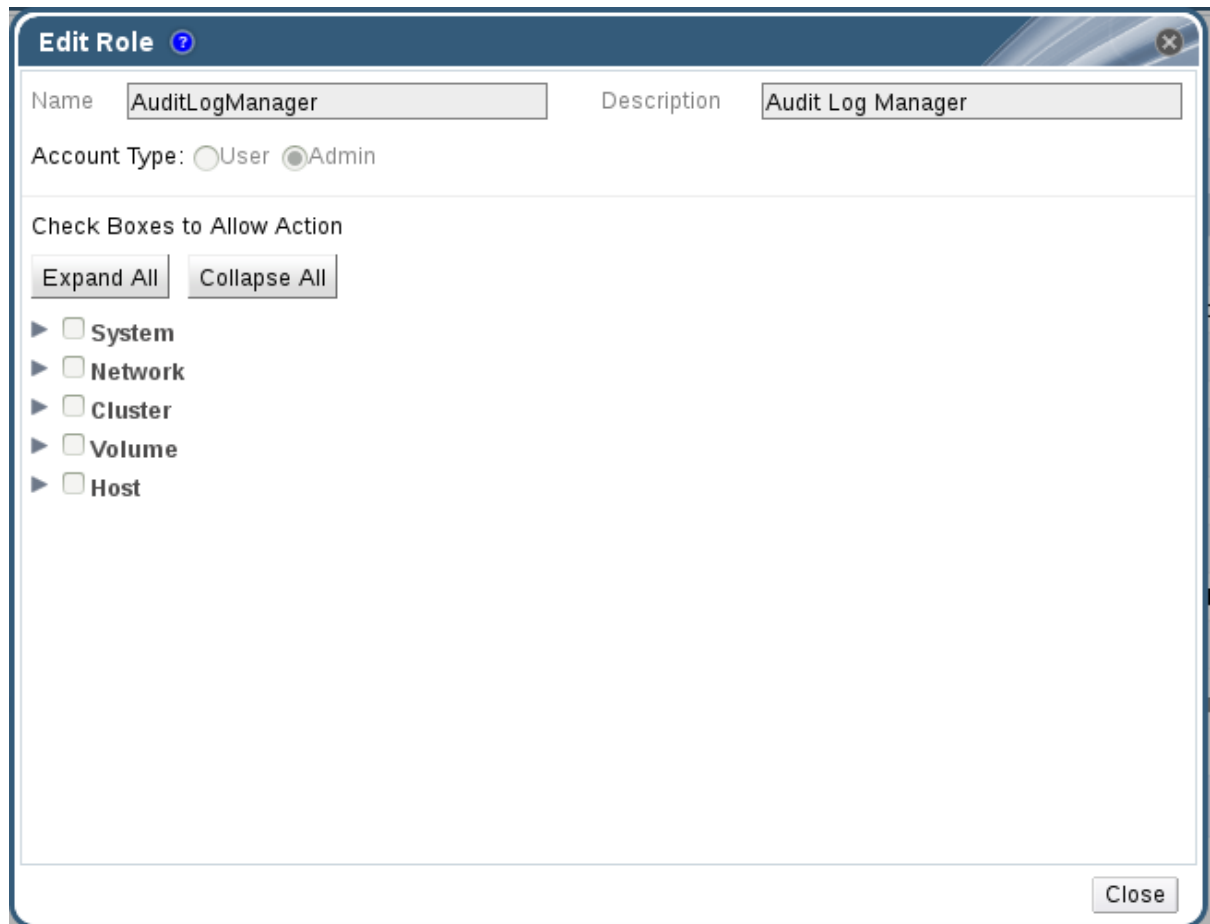


Figure 14.1. The Edit Role Dialog Box

3. If necessary, edit the **Name** and **Description** of the role. This name will display in the list of roles.
4. Use the **Expand All** or **Collapse All** buttons to view more or fewer of the permissions for the listed objects. You can also expand or collapse the options for each object.
5. For each of the objects, select or deselect the actions you wish to permit/deny for the role you are editing.
6. Click **OK** to apply the changes you have made.

14.1.4. Copying Roles

You can create a new role by cloning an existing default or custom role, and changing the permissions set as required. Use the **Copy** button on the Configure dialog box.

Procedure 14.3. Copying a Role

1. On the header bar of the Red Hat Gluster Storage Console, click **Configure**. The **Configure** dialog box displays. The dialog box includes a list of default roles, and any custom roles that exist on the Red Hat Gluster Storage Console.




Figure 14.2. The Configure Dialog Box

2. Click **Copy**. The **Copy Role** dialog box displays.
3. Change the **Name** and **Description** of the new role. This name will display in the list of roles.
4. Use the **Expand All** or **Collapse All** buttons to view more or fewer of the permissions for the listed objects. You can also expand or collapse the options for each object.
5. For each of the objects, select or deselect the actions you wish to permit/deny for the role you are editing.
6. Click **Close** to apply the changes you have made.

CHAPTER 15. BACKING UP AND RESTORING THE RED HAT GLUSTER STORAGE CONSOLE

The Red Hat Gluster Storage Console maintains important information about the environment and therefore must be regularly backed up. Regular backups ensure that Red Hat Gluster Storage Console can recover a previous state simply and quickly.

15.1. BACKING UP AND RESTORING THE RED HAT GLUSTER STORAGE CONSOLE

15.1.1. Backing up Red Hat Gluster Storage Console - Overview

While taking complete backups of the machine on which the Red Hat Gluster Storage Console is installed is recommended whenever changing the configuration of that machine, a utility is provided for backing up only the key files related to the Manager. This utility - the **engine-backup** command - can be used to rapidly back up the engine database and configuration files into a single file that can be easily stored.

15.1.2. Syntax for the engine-backup Command

The **engine-backup** command works in one of two basic modes:

```
# engine-backup --mode=backup
```

```
# engine-backup --mode=restore
```

These two modes are further extended by a set of parameters that allow you to specify the scope of the backup and different credentials for the engine database. A full list of parameters and their function is as follows:

Basic Options

--mode

Specifies whether the command will perform a backup operation or a restore operation. Two options are available - **backup**, and **restore**. This is a required parameter.

--file

Specifies the path and name of a file into which backups are to be taken in backup mode, and the path and name of a file from which to read backup data in restore mode. This is a required parameter in both backup mode and restore mode.

--log

Specifies the path and name of a file into which logs of the backup or restore operation are to be written. This parameter is required in both backup mode and restore mode.

--scope

Specifies the scope of the backup or restore operation. There are two options - **all**, which backs up both the engine database and configuration data, and **db**, which backs up only the engine database.

Database Options

--change-db-credentials

Allows you to specify alternate credentials for restoring the engine database using credentials other than those stored in the backup itself. Specifying this parameter allows you to add the following parameters.

--db-host

Specifies the IP address or fully qualified domain name of the host on which the database resides. This is a required parameter.

--db-port

Specifies the port by which a connection to the database will be made.

--db-user

Specifies the name of the user by which a connection to the database will be made. This is a required parameter.

--db-passfile

Specifies a file containing the password by which a connection to the database will be made. Either this parameter or the ***--db-password*** parameter must be specified.

--db-password

Specifies the plain text password by which a connection to the database will be made. Either this parameter or the ***--db-passfile*** parameter must be specified.

--db-name

Specifies the name of the database to which the database will be restored. This is a required parameter.

--db-secured

Specifies that the connection with the database is to be secured.

--db-secured-validation

Specifies that the connection with the host is to be validated.

Help

--help

Provides an overview of the available modes, parameters, sample usage, how to create a new database and configure the firewall in conjunction with backing up and restoring the Red Hat Gluster Storage Console.

15.1.3. Creating a Backup with the engine-backup Command

Summary

The process for creating a backup of the engine database and the configuration data for the Red Hat Gluster Storage Console using the **engine-backup** command is straightforward and can be performed while the Manager is active.

Procedure 15.1. Backing up the Red Hat Gluster Storage Console

1. Log on to the machine running the Red Hat Gluster Storage Console.
2. Run the following command to create a full backup:

Example 15.1. Creating a Full Backup

```
# engine-backup --scope=all --mode=backup --log=[file name] --
file=[file name]
```

Alternatively, run the following command to back up only the engine database:

Example 15.2. Creating an engine database Backup


```
# engine-backup --scope=db --mode=backup --log=[file name] --
file=[file name]
```

Result

A **tar** file containing a backup of the engine database, or the engine database and the configuration data for the Red Hat Gluster Storage Console, is created using the path and file name provided.

15.1.4. Restoring a Backup with the engine-backup Command

While the process for restoring a backup using the **engine-backup** command is straightforward, it involves several additional steps in comparison to that for creating a backup depending on the destination to which the backup is to be restored. For example, the **engine-backup** command can be used to restore backups to fresh installations of Red Hat Gluster Storage Console, on top of existing installations of Red Hat Gluster Storage Console, and using local or remote databases.



IMPORTANT

Backups can only be restored to environments of the same major release as that of the backup. For example, a backup of a Red Hat Gluster Storage Console version 3.4 environment can only be restored to another Red Hat Gluster Storage Console version 3.4 environment. To view the version of Red Hat Gluster Storage Console contained in a backup file, unpack the backup file and read the value in the **version** file located in the root directory of the unpacked files.

15.1.5. Restoring a Backup to a Fresh Installation

Summary

The **engine-backup** command can be used to restore a backup to a fresh installation of the Red Hat Gluster Storage Console. The following procedure must be performed on a machine on which the base operating system has been installed and the required packages for the Red Hat Gluster Storage Console have been installed, but the **engine-setup** command has not yet been run. This procedure assumes that the backup file can be accessed from the machine on which the backup is to be restored.



NOTE

The **engine-cleanup** command used to prepare a machine prior to restoring a backup only cleans the engine database, and does not drop the database, delete the user that owns that database, create engine database or perform the initial configuration of the **postgresql** service. Therefore, these tasks must be performed manually as outlined below when restoring a backup to a fresh installation.

Procedure 15.2. Restoring a Backup to a Fresh Installation

1. Log on to the machine on which the Red Hat Gluster Storage Console is installed.
2. Manually create an empty database to which the database in the backup can be restored and configure the **postgresql** service:
 1. Run the following commands to initialize the **postgresql** database, start the **postgresql** service and ensure this service starts on boot:

```
# service postgresql initdb
# service postgresql start
# chkconfig postgresql on
```

2. Run the following commands to enter the postgresql command line:

```
# su postgres
$ psql
```

3. Run the following command to create a new user:

```
postgres=# CREATE USER [user name] PASSWORD '[password]';
```

The password used while creating the database must be same as the one used while taking backup. If the password is different, follow step 3 in [Section 15.1.7, “Restoring a Backup with Different Credentials”](#)

4. Run the following command to create the new database:

```
postgres=# create database [database name] owner [user name]
template template0 encoding 'UTF8' lc_collate 'en_US.UTF-8'
lc_ctype 'en_US.UTF-8';
```

5. Edit the **/var/lib/pgsql/data/pg_hba.conf** file and add the following lines under the **'local'** section near the end of the file:

- For local databases:

```
host    [database name]    [user name]    0.0.0.0/0    md5
host    [database name]    [user name]    ::0/0        md5
```

- For remote databases:

```
host    [database name]    [user name]    X.X.X.X/32    md5
```

Replace *X.X.X.X* with the IP address of the Manager.

6. Run the following command to restart the **postgresql** service:

```
# service postgresql restart
```

3. Restore the backup using the **engine-backup** command:

```
# engine-backup --mode=restore --file=[file name] --log=[file name]
```

If successful, the following output displays:

```
Restoring...
Note: you might need to manually fix:
- iptables/firewalld configuration
- autostart of ovirt-engine service
You can now start the engine service and then restart httpd
Done.
```

4. Run the following command and follow the prompts to set up the Manager as per a fresh installation, selecting to manually configure the database when prompted:

```
# engine-setup
```

Result

The engine database and configuration files for the Red Hat Gluster Storage Console have been restored to the version in the backup.

15.1.6. Restoring a Backup to an Existing Installation

Summary

The **engine-backup** command can restore a backup to a machine on which the Red Hat Gluster Storage Console has already been installed and set up.



NOTE

The **engine-cleanup** command used to prepare a machine prior to restoring a backup only cleans the engine database, and does not drop the database, delete the user that owns that database, create engine database or perform the initial configuration of the **postgresql** service. Therefore, these tasks must be performed manually as outlined below when restoring a backup to an existing installation.

Procedure 15.3. Restoring a Backup to an Existing Installation

1. Log on to the machine on which the Red Hat Gluster Storage Console is installed.
2. Run the following command and follow the prompts to remove the configuration files for and clean the database associated with the Manager:

```
# engine-cleanup
```

Manually drop the database and create an empty database to which the database in the backup can be restored and configure the **postgresql** service

1. Run the following commands to enter the postgresql command line:

```
# su postgres
$ psql
```

2. Run the following command to drop the database:

```
# postgres=# DROP DATABASE [database name]
```

3. Run the following command to create the new database:

```
# postgres=# create database [database name] owner [user name]
template template0 encoding 'UTF8' lc_collate 'en_US.UTF-8'
lc_ctype 'en_US.UTF-8';
```

3. Restore the backup using the **engine-backup** command:

```
# engine-backup --mode=restore --file=[file name] --log=[file name]
```

If successful, the following output displays:

```
Restoring...
Note: you might need to manually fix:
- iptables/firewalld configuration
- autostart of ovirt-engine service
You can now start the engine service and then restart httpd
Done.
```

4. Run the following command and follow the prompts to re-configure the firewall and ensure the **ovirt-engine** service is correctly configured:

```
# engine-setup
```

Result

The engine database and configuration files for the Red Hat Gluster Storage Console have been restored to the version in the backup.

15.1.7. Restoring a Backup with Different Credentials

Summary

The **engine-backup** command can restore a backup to a machine on which the Red Hat Gluster Storage Console has already been installed and set up, but the credentials of the database in the backup are different to those of the database on the machine on which the backup is to be restored.



NOTE

The **engine-cleanup** command used to prepare a machine prior to restoring a backup only cleans the engine database, and does not drop the database, delete the user that owns that database, create engine database or perform the initial configuration of the **postgresql** service. Therefore, these tasks must be performed manually as outlined below when restoring a backup with different credentials.

Procedure 15.4. Restoring a Backup with Different Credentials

1. Log on to the machine on which the Red Hat Gluster Storage Console is installed.
2. Run the following command and follow the prompts to remove the configuration files for and clean the database associated with the Manager:

```
# engine-cleanup
```

Manually drop the database and create an empty database to which the database in the backup can be restored and configure the **postgresql** service:

1. Run the following commands to enter the postgresql command line:

```
# su postgres
$ psql
```

2. Run the following command to drop the database:

```
# postgres=# DROP DATABASE [database name]
```

3. Run the following command to create the new database:

```
# postgres=# create database [database name] owner [user name]
template template0 encoding 'UTF8' lc_collate 'en_US.UTF-8'
lc_ctype 'en_US.UTF-8';
```

3. Restore the backup using the **engine-backup** command with the **--change-db-credentials** parameter:

```
# engine-backup --mode=restore --file=[file name] --log=[file name]
--change-db-credentials --db-host=[database location] --db-
name=[database name] --db-user=[user name] --db-password=[password]
```

If successful, the following output displays:

```
Restoring...
Note: you might need to manually fix:
- iptables/firewalld configuration
- autostart of ovirt-engine service
You can now start the engine service and then restart httpd
Done.
```

4. Run the following command and follow the prompts to re-configure the firewall and ensure the **ovirt-engine** service is correctly configured:

```
█ # engine-setup
```

Result

The engine database and configuration files for the Red Hat Gluster Storage Console have been restored to the version in the backup using the supplied credentials.

APPENDIX A. UTILITIES

A.1. DOMAIN MANAGEMENT TOOL

Red Hat Gluster Storage Console uses directory services to authenticate users. During installation, the manager sets up a domain named **internal** that is only used to store the **admin** user. To add and remove other users from the system, you must first add the directory services in which they are found.

The supported directory service is IPA. Red Hat Gluster Storage Console includes a domain management tool, **rhsc-manage-domains**, to add and remove domains provided by this service. In this way, you can grant access to the Red Hat Gluster Storage environment to users stored across multiple domains.

You will find the **rhsc-manage-domains** command on the machine on which Red Hat Gluster Storage Console was installed. The **rhsc-manage-domains** command must be run as the root user.

A.1.1. Syntax

The usage syntax is:

```
# rhsc-manage-domains action [options]
```

The available actions are:

add

Add a domain to the console directory services configuration.

edit

Edit a domain in the console directory services configuration.

delete

Delete a domain from the console directory services configuration.

validate

Validate the console directory services configuration. The command attempts to authenticate to each domain in the configuration using the configured user name and password.

list

List the current directory services configuration of the console.

The options that can be combined with the actions on the command line are:

--domain=DOMAIN

Specifies the domain on which the action must be performed. The **--domain** parameter is mandatory for **add**, **edit**, and **delete**.

--user=USER

Specifies the domain user to use. The **--user** parameter is mandatory for **add**, and optional for **edit**.

--password-file=FILE

A file containing the password. If this is not set, the password is read interactively.

--config-file=FILE

Specifies an alternative configuration file that the command must load. The **--config-file** parameter is always optional.

--report

Specifies that all validation errors encountered while performing the validate action will be reported in full.

Common examples of usage are discussed in subsequent sections. For full information on usage, see the **rhsc-manage-domains** command help output:

```
# rhsc-manage-domains --help
```

A.1.2. Listing Domains in Configuration

The `rhsc-manage-domains` command lists the directory services domains defined in the Red Hat Gluster Storage Console configuration. This command prints the domain, the user name in User Principal Name (UPN) format, and whether the domain is local or remote for each configuration entry.

Example A.1. rhsc-manage-domains List Action

```
# rhsc-manage-domains list
Domain: directory.demo.redhat.com
User name: admin@DIRECTORY.DEMO.REDHAT.COM
This domain is a remote domain.
```

A.1.3. Adding Domains to Configuration

In this example, the `rhsc-manage-domains` command is used to add the IdM domain `directory.demo.redhat.com` to the Red Hat Gluster Storage Console configuration. The configuration is set to use the `admin` user when querying the domain; the password is provided interactively.

Example A.2. rhsc-manage-domains Add Action

```
# rhsc-manage-domains add --domain=directory.demo.redhat.com --
provider=IPA --user=admin
loaded template kr5.conf file
setting default_tkt_enctypes
setting realms
setting domain realm
success
User guid is: 80b71bae-98a1-11e0-8f20-525400866c73
Successfully added domain directory.demo.redhat.com. oVirt Engine
restart is required in order for the changes to take place (service
ovirt-engine restart).
```

A.1.4. Editing a Domain in the Configuration

In this example, the `rhsc-manage-domains` command is used to edit the `directory.demo.redhat.com` domain in the Red Hat Gluster Storage Console configuration. The configuration is updated to use the `admin` user when querying this domain; the password is provided interactively.

Example A.3. `rhsc-manage-domains` Edit Action

```
# rhsc-manage-domains edit --domain=directory.demo.redhat.com --
user=admin
loaded template kr5.conf file
setting default_tkt_enctypes
setting realms
setting domain realm0
success
User guide is: 80b71bae-98a1-11e0-8f20-525400866c73
Successfully edited domain directory.demo.redhat.com. oVirt Engine
restart is required in order for the changes to take place (service
ovirt-engine restart).
```

A.1.5. Validating Domain Configuration

In this example, the `rhsc-manage-domains` command is used to validate the Red Hat Gluster Storage Console configuration. The command attempts to log into each listed domain with the credentials provided in the configuration. The domain is reported as valid if the attempt is successful.

Example A.4. `rhsc-manage-domains` Validate Action

```
# rhsc-manage-domains validate
User guide is: 80b71bae-98a1-11e0-8f20-525400866c73
Domain directory.demo.redhat.com is valid.
```

A.1.6. Deleting a Domain from the Configuration

In this example, the `rhsc-manage-domains` command is used to remove the `directory.demo.redhat.com` domain from the Red Hat Gluster Storage Console configuration. Users defined in the removed domain will no longer be able to authenticate with the Red Hat Gluster Storage Console. The entries for the affected users will remain defined in the Red Hat Gluster Storage Console until they are explicitly removed.

The domain being removed in this example is the last one listed in the Red Hat Gluster Storage Console configuration. A warning is displayed highlighting this fact and that only the `admin` user from the internal domain will be able to log in until another domain is added.

Example A.5. `rhsc-manage-domains` Delete Action

```
# rhsc-manage-domains delete --domain=directory.demo.redhat.com
WARNING: Domain directory.demo.redhat.com is the last domain in the
configuration. After deleting it you will have to either add another
domain, or to use the internal admin user in order to login.
```

Successfully deleted domain `directory.demo.redhat.com`. Please remove all users and groups of this domain using the Administration portal or the API.

APPENDIX B. CHANGING PASSWORDS IN RED HAT GLUSTER STORAGE CONSOLE

This appendix describes how to change passwords for the administrator user in the Administration Portal and Red Hat Gluster Storage Console PostgreSQL databases.

B.1. CHANGING THE PASSWORD FOR THE ADMINISTRATOR USER

The **admin@internal** user account is automatically created on installing and configuring Red Hat Gluster Storage Console. This account is stored locally in the Red Hat Gluster Storage Console PostgreSQL database and exists separately from other directory services. Unlike IPA domains, users cannot be added to or deleted from the internal domain. The **admin@internal** user is the SuperUser for Red Hat Gluster Storage Console, and has administrator privileges over the environment via the Administration Portal.

During installation, you were prompted to set a password for the **admin@internal** user. However, if you have forgotten the password or choose to reset the password, you can use the **rhsc-config** utility.

Procedure B.1. Resetting the Password for the admin@internal User

1. Log in to the Red Hat Gluster Storage Console server as the **root** user.
2. Use the **rhsc-config** utility to set a new password for the **admin@internal** user. Run the following command:

```
# rhsc-config -s AdminPassword=interactive
```

After typing the above command, a password prompt displays for you to enter the new password.

You do not need to use quotes. However, use escape shell characters if you include them in the password.

3. Restart the **ovirt-engine** service to apply the changes. Run the following command:

```
# service ovirt-engine restart
```

APPENDIX C. SEARCH PARAMETERS

This appendix describes the search function in Red Hat Gluster Storage Console in detail.

C.1. SEARCH QUERY SYNTAX

Each part of the query syntax is explained in greater detail below.

Example	Result
Hosts: cluster = cluster name	Displays a list of all servers in the cluster.
Volumes: status = up	Displays a list of all volumes with status up.
Events: severity > normal sortby time	Displays the list of all events whose severity is higher than Normal , sorted by time.

As you type each part of a search query, a drop-down list of choices for the next part of the search opens below the **Search** bar. You can either select from the list and then continue typing or selecting the next part of the search, or ignore the options and continue entering your query manually.

The following table shows how Red Hat Gluster Storage Console auto-completion assists in query construction. It shows what the drop-down list will display as the administrator inputs text into the search field.

Hosts: cluster = down

Input	List Items Displayed	Action
h	Hosts (1 option only)	Select Hosts or; Type Hosts
Hosts:	All host properties	Type c
Hosts: c	host properties starting with c	Select cluster or type cluster
Hosts: cluster	= =!	Select or type =
Hosts: cluster =		Select or type cluster name

C.2. SEARCHING FOR RESOURCES

This section specifies the unique set of properties for each resource and the set of associated resource types.

C.2.1. Searching for Clusters

The following table describes all search options for clusters.

Property (of resource or resource-type)	Type	Description (Reference)
name	String	The unique name that identifies the clusters on the network.
description	String	The description of the cluster.
initialized	String	A Boolean True or False indicating the status of the cluster.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Clusters: initialized = true or name = Default

This above query returns a list of clusters that are:

- Initialized; or
- Named **Default**

C.2.2. Searching for Hosts

The following table describes all search options for hosts.

Property (of resource or resource-type)	Type	Description (Reference)
Events.events-prop	See property types in Section C.2.5, "Searching for Events"	The property of the events associated with the host.
Users.users-prop	See property types in Section C.2.4, "Searching for Users"	The property of the users associated with the host.
name	String	The name of the host.
status	List	The availability of the host.

Property (of resource or resource-type)	Type	Description (Reference)
cluster	String	The cluster to which the host belongs.
address	String	The unique name that identifies the host on the network.
cpu_usage	Integer	The percent of processing power usage.
mem_usage	Integer	The percentage of memory usage.
network_usage	Integer	The percentage of network usage.
load	Integer	Jobs waiting to be executed in the <i>run-queue</i> per processor, in a given time slice.
version	Integer	The version number of the operating system.
cpus	Integer	The number of CPUs on the host.
memory	Integer	The amount of memory available.
cpu_speed	Integer	The processing speed of the CPU.
cpu_model	String	The type of CPU.
committed_mem	Integer	The percentage of committed memory.
tag	String	The tag assigned to the host.
type	String	The type of host.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Host: cluster = Default

The above query returns a list of hosts that:

- Are part of the **Default** cluster.

C.2.3. Searching for Volumes

The following table describes all search options for volumes.

Property (of resource or resource-type)	Type	Description (Reference)
Clusters.clusters prop	See property types in Section C.2.1, “Searching for Clusters”	The property of the clusters associated with the volume.
name	String	The name of the volume.
status	List	The availability of the volume.
type	List	The type of the volume.
transport_type	List	The transport type of the volume.
replica_count	Integer	The replicate count of the volume.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Volumes: Cluster.name = Default and Status = Up

The above query returns a list of volumes that:

- Belong to the **Default** cluster and the status of the volume is **Up**.

C.2.4. Searching for Users

The following table describes all search options for users.

Property (of resource or resource-type)	Type	Description (Reference)
Hosts.hosts- prop	See property types in Section C.2.2, “Searching for Hosts”	The property of the hosts associated with the user.

Property (of resource or resource-type)	Type	Description (Reference)
Events.events-prop	See property types in Section C.2.5, "Searching for Events"	The property of the events associated with the user.
name	String	The name of the user.
lastname	String	The last name of the user.
username	String	The unique name of the user.
department	String	The department to which the user belongs.
group	String	The group to which the user belongs.
title	String	The title of the user.
status	String	The status of the user.
role	String	The role of the user.
tag	String	The tag to which the user belongs.
pool	String	The pool to which the user belongs.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Users: Events.severity > normal and Hosts.name = Server name

The above query returns a list of users for which:

- Events of a severity greater than **Normal** have occurred on their hosts.

C.2.5. Searching for Events

The following table describes all search options you can use to search for events. Auto-completion is offered for many options as appropriate.

Property (of resource or resource-type)	Type	Description (Reference)
Hosts.hosts-prop	See property types in Section C.2.2, “Searching for Hosts”	The property of the hosts associated with the event.
Users.users-prop	See property types in Section C.2.4, “Searching for Users”	The property of the users associated with the event.
type	List	Type of the event.
severity	List	The severity of the Event: Warning/Error/Normal
message	String	Description of the event type.
time	Integer	Time at which the event occurred.
username	username	The user name associated with the event.
event_host	String	The host associated with the event.
sortby	List	Sorts the returned results by one of the resource properties.
page	Integer	The page number of results to display.

Example

Events: `event_host = gonzo.example.com`

The above query returns a list of events for which:

- The event occurred on the server named `gonzo.example.com`.

C.3. SAVING AND ACCESSING QUERIES AS BOOKMARKS

Search queries can be saved as bookmarks. This allows you to sort and display results lists with a single click. You can save, edit and remove bookmarks with the **Bookmarks** pane.

C.3.1. Creating Bookmarks

Bookmarks can be created for any type of available search, using a number of criteria.

Procedure C.1. Saving a Query String as a Bookmark

1. Enter the search query in the **Search** bar (see Appendix D).
2. Click the **Bookmark** button to the right of the **Search** bar.

The **New Bookmark** dialog box displays. The query displays in the **Search String** field. You can edit the query if required.

3. In **Name**, specify a descriptive name for the search query.
4. Click **OK** to save the query as a bookmark.
5. The search query is saved and displays in the **Bookmarks** pane.

C.3.2. Editing Bookmarks

Bookmarks can be edited for any type of available search, using an existing bookmark.

Procedure C.2. Editing a Bookmark

1. Select a bookmark from the **Bookmarks** pane.
2. The results list displays the items according to the criteria. Click the **Edit** button on the **Bookmark** pane.

The **Edit Bookmark** dialog box displays. The query displays in the **Search String** field. Edit the search string as required.

3. Change the **Name** and **Search String** as necessary.
4. Click **OK** to save the edited bookmark.

C.3.3. Deleting Bookmarks

Bookmarks can be deleted.

Procedure C.3. Deleting a Bookmark

1. Select one or more bookmark from the **Bookmarks** pane.
2. The results list displays the items according to the criteria. Click the **Remove** button on the **Bookmark** pane.

The **Remove Bookmark** dialog box displays.

3. Click **OK** to remove the selected bookmarks.

APPENDIX D. CONFIGURATION FILES

D.1. NAGIOS CONFIGURATION FILES

Auto-discovery creates folders and files as part of configuring Red Hat Gluster Storage nodes for monitoring. All nodes in the trusted storage pool are configured as hosts in Nagios. The Host and Hostgroup configurations are also generated for trusted storage pool with cluster name. Ensure that the following files and folders are created with the details described to verify the Nagios configurations generated using Auto-discovery.

- In `/etc/nagios/gluster/` directory, a new directory **Cluster-Name** is created with the name provided as **Cluster-Name** while executing **configure-gluster-nagios** command for auto-discovery. All configurations created by auto-discovery for the cluster are added in this folder.
- In `/etc/nagios/gluster/Cluster-Name` directory, a configuration file, **Cluster-Name.cfg** is generated. This file has the host and hostgroup configurations for the cluster. This also contains service configuration for all the cluster/volume level services.

The following Nagios object definitions are generated in **Cluster-Name.cfg** file:

- A hostgroup configuration with **hostgroup_name** as cluster name.
- A host configuration with **host_name** as cluster name.
- The following service configurations are generated for cluster monitoring:
 - A *Cluster - Quorum* service to monitor the cluster quorum.
 - A *Cluster Utilization* service to monitor overall utilization of volumes in the cluster. This is created only if there is any volume present in the cluster.
 - A *Cluster Auto Config* service to periodically synchronize the configurations in Nagios with Red Hat Gluster Storage trusted storage pool.
- The following service configurations are generated for each volume in the trusted storage pool:
 - A *Volume Status - Volume-Name* service to monitor the status of the volume.
 - A *Volume Utilization - Volume-Name* service to monitor the utilization statistics of the volume.
 - A *Volume Quota - Volume-Name* service to monitor the Quota status of the volume, if Quota is enabled for the volume.
 - A *Volume Self-Heal - Volume-Name* service to monitor the Self-Heal status of the volume, if the volume is of type replicate or distributed-replicate.
 - A *Volume Geo-Replication - Volume-Name* service to monitor the Geo Replication status of the volume, if Geo-replication is configured for the volume.
- In `/etc/nagios/gluster/Cluster-Name` directory, a configuration file with name **Host-Name.cfg** is generated for each node in the cluster. This file has the host configuration for the node and service configuration for bricks from the particular node. The following Nagios object definitions are generated in **Host-name.cfg**.

- A host configuration which has *Cluster-Name* in the **hostgroups** field.
- The following services are created for each brick in the node:
 - A *Brick Utilization - brick-path* service to monitor the utilization of the brick.
 - A *Brick - brick-path* service to monitor the brick status.

Table D.1. Nagios Configuration Files

File Name	Description
<code>/etc/nagios/nagios.cfg</code>	Main Nagios configuration file.
<code>/etc/nagios/cgi.cfg</code>	CGI configuration file.
<code>/etc/httpd/conf.d/nagios.conf</code>	Nagios configuration for httpd.
<code>/etc/nagios/passwd</code>	Password file for Nagios users.
<code>/etc/nagios/nrpe.cfg</code>	NRPE configuration file.
<code>/etc/nagios/gluster/gluster-contacts.cfg</code>	Email notification configuration file.
<code>/etc/nagios/gluster/gluster-host-services.cfg</code>	Services configuration file that's applied to every Red Hat Gluster Storage node.
<code>/etc/nagios/gluster/gluster-host-groups.cfg</code>	Host group templates for a Red Hat Gluster Storage trusted storage pool.
<code>/etc/nagios/gluster/gluster-commands.cfg</code>	Command definitions file for Red Hat Gluster Storage Monitoring related commands.
<code>/etc/nagios/gluster/gluster-templates.cfg</code>	Template definitions for Red Hat Gluster Storage hosts and services.
<code>/etc/nagios/gluster/snmpmanagers.conf</code>	SNMP notification configuration file with the IP address and community name of SNMP managers where traps need to be sent.

APPENDIX E. REVISION HISTORY

Revision 3.4-1

Tue Sep 04 2018

**Red Hat Gluster Storage
Documentation Team**

Updates for Red Hat Gluster Storage 3.4 release.