
Gluster Management Console - 1.0.0

User Guide

Table of Contents

| | | |
|--------|---|----|
| 1. | About the Guide | 4 |
| 1.1. | Disclaimer | 4 |
| 1.2. | Audience | 4 |
| 1.3. | Terminology | 4 |
| 1.4. | Conventions | 4 |
| 1.5. | Documentation Feedback..... | 5 |
| 2. | Introducing Gluster Management Console | 6 |
| 2.1. | Gluster FileSystem | 6 |
| 2.2. | Gluster Management Gateway | 7 |
| 2.3. | Gluster Management Console | 7 |
| 3. | Gluster Management Console License | 8 |
| 4. | Installing Gluster Management Gateway..... | 9 |
| 4.1. | Checking the Gluster Management Gateway Minimum Requirements | 9 |
| 4.2. | Installing the Gluster Management Gateway | 9 |
| 4.3. | Installing Gluster Management Gateway Backend Package..... | 10 |
| 5. | Setting up the Gluster Management Console..... | 12 |
| 5.1. | Checking the Minimum Requirements on the Client Machine | 12 |
| 5.2. | Accessing the Gluster Management Console Using Java Web Start | 12 |
| 6. | Exploring the Gluster Management Console User Interface | 15 |
| 6.1. | Gluster Management Console Toolbar | 16 |
| 6.2. | Tasks..... | 17 |
| 6.3. | Terminal | 17 |
| 6.4. | Monitoring Cluster Resources | 19 |
| 6.5. | Alerts | 19 |
| 6.6. | Preferences | 20 |
| 6.6.1. | Setting Gluster Preferences | 20 |
| 6.6.2. | Setting Alerts Preferences | 21 |
| 6.6.3. | Setting Charts Preferences | 22 |
| 6.6.4. | Setting Terminal Preferences | 23 |
| 6.7. | Security Keys..... | 24 |
| 6.7.1. | Importing Keys | 24 |
| 6.7.2. | Exporting Keys | 24 |
| 6.8. | Resetting the Password | 24 |
| 6.9. | CIFS Services | 24 |
| 6.9.1. | Configuring the CIFS Services | 25 |

| | | |
|--------|---|----|
| 6.9.2. | Creating CIFS Users..... | 25 |
| 6.9.3. | Removing CIFS Users | 25 |
| 7. | Managing Clusters | 26 |
| 7.1. | Creating a Cluster | 27 |
| 7.2. | Registering a Cluster | 27 |
| 7.3. | Selecting a Cluster | 28 |
| 8. | Managing Servers | 29 |
| 8.1. | Adding Servers to a Cluster..... | 30 |
| 8.2. | Removing Servers from a Cluster..... | 31 |
| 8.3. | Initializing a Disk | 32 |
| 9. | Managing Volumes..... | 34 |
| 9.1. | Creating a Volume | 35 |
| 9.2. | Starting Volumes | 36 |
| 9.3. | Tuning Volume Options | 37 |
| 9.4. | Stopping Volumes | 38 |
| 9.5. | Rebalancing Volume | 39 |
| 9.6. | Deleting Volumes..... | 40 |
| 9.7. | Managing Bricks..... | 41 |
| 9.7.1. | Adding Bricks | 41 |
| 9.7.2. | Removing Bricks | 42 |
| 9.7.3. | Migrating Bricks..... | 43 |
| 10. | Managing Logs..... | 45 |
| 10.1. | Displaying Log Based on a Search Criteria | 45 |
| 10.2. | Downloading the Logs | 46 |

1. About the Guide

This guide describes the basic overview and installation procedures for the Gluster Management Console. It also describes how to use the Gluster Management Console to build, manage, and maintain your Gluster storage cloud.

1.1. Disclaimer

Gluster, Inc. has designated English as the official language for all of its product documentation and other documentation, as well as all our customer communications. All documentation prepared or delivered by Gluster will be written, interpreted and applied in English, and English is the official and controlling language for all our documents, agreements, instruments, notices, disclosures and communications, in any form, electronic or otherwise (collectively, the “Gluster Documents”).

Any customer, vendor, partner or other party who requires a translation of any of the Gluster Documents is responsible for preparing or obtaining such translation, including associated costs. However, regardless of any such translation, the English language version of any of the Gluster Documents prepared or delivered by Gluster shall control for any interpretation, enforcement, application or resolution.

1.2. Audience

This book is intended for anyone who wants to configure and manage a Gluster cloud storage environment. The information in this book is written for experienced system administrators who are familiar with the Linux operating system, concepts of File System, and GlusterFS concepts.

1.3. Terminology

The terms defined here help you understand the concepts in this document.

| Term | Description |
|---------|---|
| Brick | Brick is the basic unit of storage, represented by an export directory on a server in the trusted storage pool. |
| Cluster | A cluster is a group of linked computers, working together closely thus in many respects forming a single computer. |
| Volume | A volume is a logical collection of bricks. Most of the Gluster management operations happen on the volume. |

1.4. Conventions

The following table lists the formatting conventions that are used in this guide to make it easier for you to recognize and use specific types of information.

| Convention | Description | Example |
|--------------|--|---|
| Courier Text | Commands formatted as courier indicate shell commands. | <code>gluster volume start volname</code> |

| Convention | Description | Example |
|------------------------|--|---|
| <i>Italicized Text</i> | Within a command, italicized text represents variables, which must be substituted with specific values. | <code>gluster volume start <i>volname</i></code> |
| Square Brackets | Within a command, optional parameters are shown in square brackets. | <code>gluster volume start volname [force]</code> |
| Curly Brackets | Within a command, alternative parameters are grouped within curly brackets and separated by the vertical OR bar. | <code>gluster volume { start stop delete } volname</code> |

1.5. Documentation Feedback

Gluster welcomes your comments and suggestions on the quality and usefulness of its documentation. If you find any errors or have any other suggestions, write to us at docfeedback@gluster.com for clarification. Provide the chapter, section, and page number, if available.

Gluster offers a range of resources related to Gluster software:

- Discuss technical problems and solutions on the Discussion Forum (<http://community.gluster.org>)
- Get hands-on step-by-step tutorials (http://www.gluster.com/community/documentation/index.php/Main_Page)
- Reach Support (<http://www.gluster.com/services/>)

2. Introducing Gluster Management Console

This section introduces the Gluster Filesystem, Gluster Management Gateway, and Gluster Management Console. It also introduces how to use the Gluster Management Console to manage, optimize, and troubleshoot your storage environment.

2.1. Gluster FileSystem


GlusterFS is an open source, scale-out file system capable of scaling to several petabytes and handling thousands of clients. GlusterFS assembles storage building blocks, aggregating disk and memory resources and managing data in a single global namespace. GlusterFS is based on a stackable user space design and can deliver exceptional performance for diverse workloads.

The GlusterFS modular architecture allows administrators to stack modules to match user requirements, as needed. For example, administrators can use GlusterFS to quickly configure a standalone server system and later expand the system as their business needs grow.

Virtualized Cloud Environments

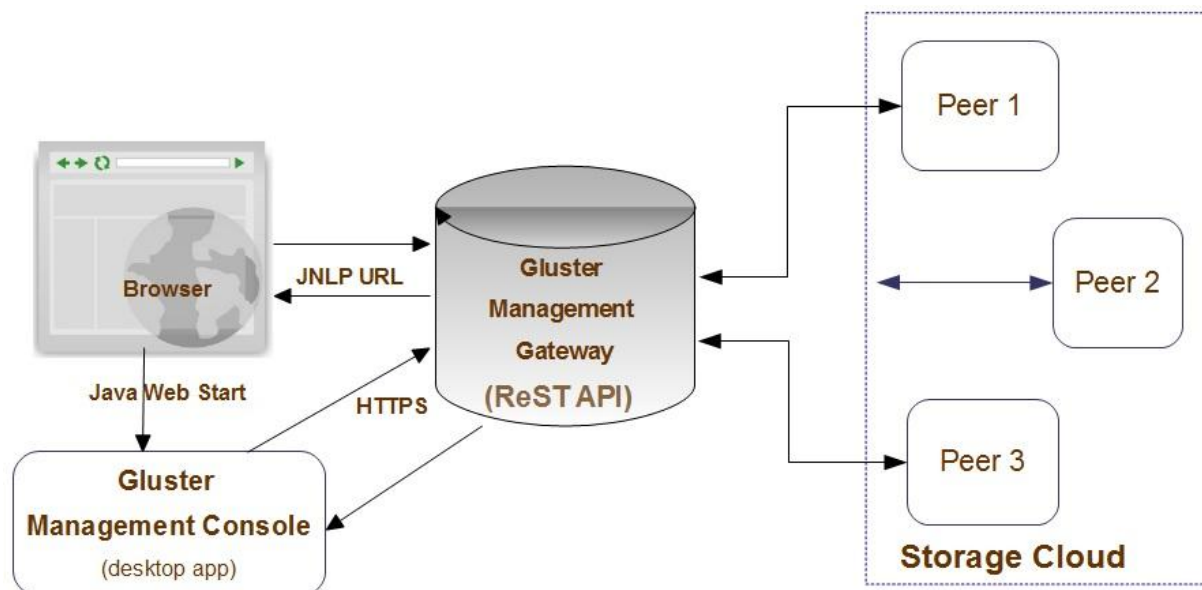
GlusterFS is designed for today's high-performance, virtualized cloud environments. Unlike traditional data centers, cloud environments require multi-tenancy along with the ability to grow or shrink resources on demand. GlusterFS incorporates cloud capabilities into the core architecture.

Every module in GlusterFS is treated as a volume. GlusterFS further adds an elastic volume manager (the `glusterd` daemon) and a console manager (the `gluster` command line interface). Using the console manager, administrators can dynamically expand, shrink, rebalance, and migrate volumes, among other operations. The command line interface additionally provides an interactive shell (with auto-completion) along with scripting capabilities suitable for automation.



2.2. Gluster Management Gateway

Gluster Management Gateway provides simple and powerful ReST APIs for managing your Gluster Storage Cluster. It also hosts the Gluster Management Console binaries, which can be launched from a browser using the Java Webstart. Gluster Management Console uses the ReST APIs for performing the administrative tasks on the Gluster Storage Cluster.



2.3. Gluster Management Console

Gluster Management Console is a Java application that provides a single interface to perform the administrative tasks such as create, manage, and maintain your storage cluster.

You can use the Gluster Management Console to:

- Create a Gluster storage cluster
- Expand your cluster
- Monitor systems in your cluster
- Manage logs

3. Gluster Management Console License

The Gluster Management Console End User License Agreement (EULA) is available at <http://www.gluster.com/eula/>.

4. Installing Gluster Management Gateway

Your network storage environment must have the required software, hardware, disk space, and memory to install the Gluster Management Gateway. The installation procedure terminates if it detects that your environment does not meet the minimum installation requirements. You must ensure to configure the missing requirements before you can attempt to install the software again.

This chapter provides an overview of the system requirements for installing the Gluster Management Gateway.

For checking the GlusterFS minimum requirements, refer to the [Gluster 3.2 Filesystem Installation Guide, Checking Minimum Requirements](#).

4.1. Checking the Gluster Management Gateway Minimum Requirements

Gluster Management Gateway is installed on a central server using the Linux package published by Gluster. Before you install Gluster Management Gateway, you must verify that your environment matches the minimum requirements described in this section.

You must verify that your central server matches the minimum requirements listed below:

- Memory: 2 GB RAM
- Processor: 2 (2.4 GHz each) is recommended; 64 bit
- Disk Space: 10 GB
- The gateway machine can be either Gluster Storage Software Appliance, Gluster Virtual Storage Appliance, or CentOS 5.6.

4.2. Installing the Gluster Management Gateway

To install the Gluster Management Gateway, perform the following steps:

1. Download the Gluster Management Gateway repository configuration file (*glustermg.repo*) and the Gluster Management Gateway RPM (*glustermg-1.0.0-1.x86_64.rpm*) from the location provided by Gluster.

Note: *glustermg.repo* has been created and tested with CentOS 5.6. It may work with other versions of CentOS, however it has not been extensively tested.

2. Save *glustermg.repo* in the following location on your machine:

```
/etc/yum.repos.d/
```

3. Run the following command to install *glustermg* dependencies:

```
# yum groupinstall glustermg
```

4. The following packages are installed if not installed earlier:

- Java SE runtime version 1.6 - OpenJDK v1.6.0.0
- Tomcat v5.5.23
- wget

5. Run the following command to install Gluster Management Gateway.

```
rpm -ivh glustermg-1.0.0-1.x86_64.rpm
```

Gateway generates SSH keys during installation, which are stored in `/opt/glustermg/keys`. These keys are used for communication with the storage servers. However, in Amazon Web Services (AWS) environment, SSH communication is done using the private key provided by Amazon. Hence, the AWS generated private key (`.pem` file) must be stored in `/opt/glustermg/keys` with filename as `gluster.pem`

Note:

- You can install the Gluster Management Gateway packages on a Gluster Storage Server. However, this can impact performance of the GlusterFS processes on this machine hence this is not recommended.
- The Gluster Management Gateway is shipped with a self-signed SSL certificate, by default. While this ensures that the communication between the console and gateway is encrypted and safe, some browsers will issue warning when you access the server. It is highly recommended that you obtain your own service certificate from a valid signing authority, install and configure Tomcat5 for the 8443 port. For more information on obtaining your own service certificate, refer to [Installing a Certificate from a Certificate Authority](#) from Tomcat documentation.
- Ensure to enable port 8443 (HTTPS), port 24731 (TCP), and port 24729 (UDP) on the gateway.
- Gluster Management Gateway communicates with all storage servers through SSH, using their host names. Make sure all of your storage server's host names (fully qualified domain names) are resolvable by the Gateway and by each other. If you have a properly configured DNS setup, this should work seamlessly.

4.3. Installing Gluster Management Gateway Backend Package

Gluster Management Gateway communicates with every node of your cluster using the backend package. This package must be installed on each and every node of your cluster using the Linux package published by Gluster.

Pre-requisite

- Your backend machines can be either Gluster Storage Software Appliance, Gluster Virtual Storage Appliance, or CentOS 5.6 running GlusterFS v3.2.3 or higher.

To install the backend packages on each and every node of your cluster, perform the following steps:

1. Download the Gluster Management Gateway repository configuration file (`glustermg.repo`) and the Gluster Management Gateway Backend RPM (`glustermg-backend-1.0.0-1.x86_64.rpm`) from the location provided by Gluster.

Note: *glustermg.repo* has been created and tested with CentOS 5.6. It may work with other versions of CentOS, however it has not been extensively tested.

2. Save *glustermg.repo* in the following location on your machine:

```
/etc/yum.repos.d/
```

3. Run the following command:

```
# yum groupinstall glustermg-backend
```

4. Run the following command to install the *glustermg-backend* dependencies:

```
rpm -ivh glustermg-backend-1.0.0-1.x86_64.rpm
```

5. The following backend packages are installed, if not installed earlier:

- chkconfig
- python v2.4.3
- perl v5.8.8
- rrdtool-perl v1.2.27
- appliance-base v1.3 (published by Gluster)
- sudo
- crontabs
- samba3
- libxml2 v2.6

Note:

- You must enable port 24729 (UDP) for auto discovery of servers.
- In case CIFS only, the Samba services must run on port 445 which need to be enabled and open on all the storage servers.
- Paravirtual VMs on Xen are not supported.

5. Setting up the Gluster Management Console

The section describes how you can set up the Gluster Management Console after installing the Gluster Management Gateway.

5.1. Checking the Minimum Requirements on the Client Machine

Verify the following:

- Processor - 1.4 GHz
- Memory - 2 GB RAM
- Disk - 100 MB
- Operating Systems: Windows XP, Vista, or Windows 7, Linux, Mac OS X
- Oracle JRE 1.6

Web Browser Requirements

- Use your Java Runtime Environment (JRE) installer to configure your browser for Java Web Start. Also, ensure that you have configured your browser to open `.jnlp` files using the Java Web Start application or `javaws.exe` located in `bin` folder of your JRE installation.
- When you access the Gluster Management Console, Java Web Start downloads all required `.jar` files from the Gateway server and launches the console's GUI. During this process, web start may display a certificate warning mentioning that the application is signed by Gluster, Inc. This certificate is authentic, but it requires full permissions as the application contains features that require registering information on the client's hard disk such as exporting the security keys from gateway. You can accept the warning and select **Always Trust Content from this Provider** on your browser to avoid this warning every time you launch the application.

5.2. Accessing the Gluster Management Console Using Java Web Start

To start Gluster Management Console using Java web start, perform the following steps.

1. Open your web browser.

Ensure that your web browser has cookies and Javascript enabled.

2. Type the following URL in your web browser:

```
https://server-name:8443/glustermg/
```

Gluster Management Console login window is displayed.



3. Type the following default user name and password in the login window:

Username: gluster

Password: syst3m

Note: In Mac OS, it is not possible to identify the CPU architecture through the Java script. The system provides links to all possible combinations, click on the appropriate link to launch the application.

4. When you access the application for the first time, the console requests you to change this password.



5. Type your current password in the **Old Password** field.
The default system-assigned password is *syst3m*
6. Type your new password in the **New Password** field.
7. Retype your new password in the **Confirm Password** field.
8. Click **Change**. The system displays **Change Password** window indicating that the password was successfully changed.



9. Click **OK**. The Select Cluster window is displayed. You can now create a new cluster or register an existing cluster. For more information on creating or registering clusters, refer to [Chapter 6, Managing Clusters](#).

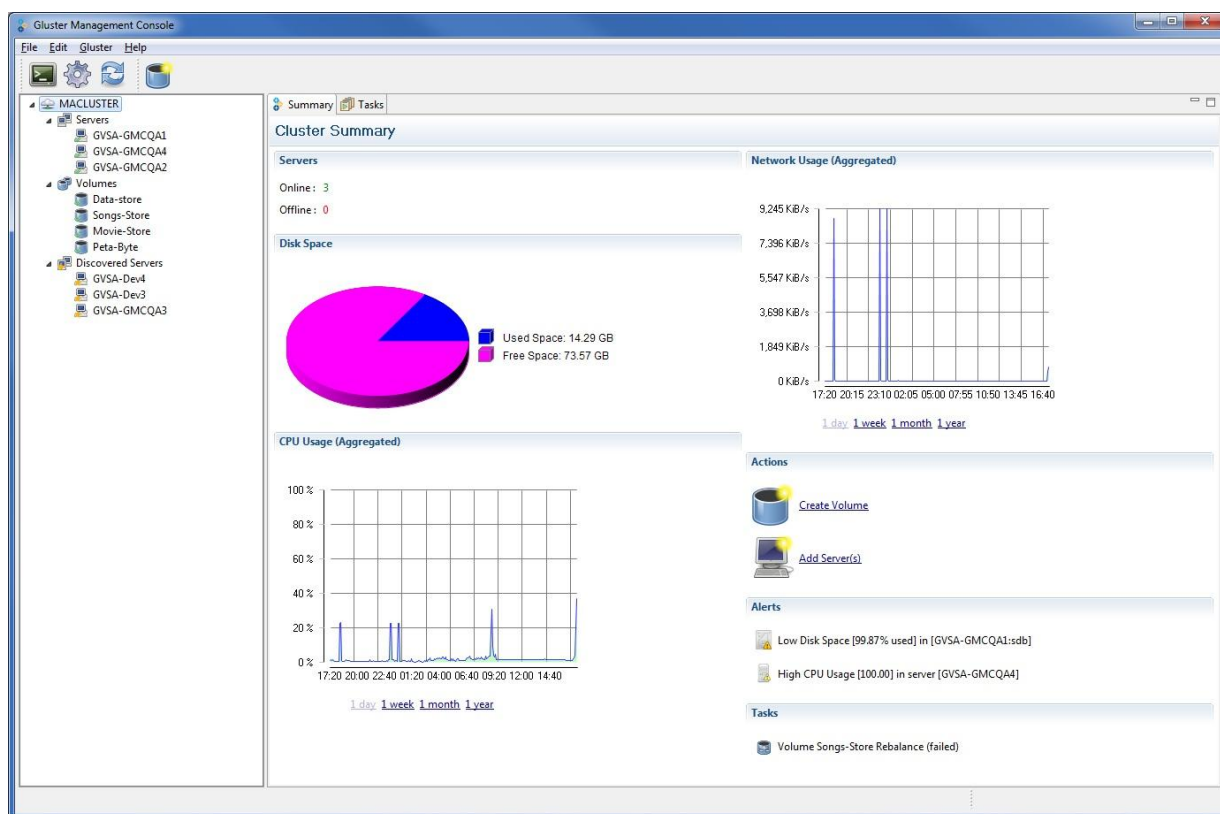
Note: You can change the password at any time by accessing **Edit > Change Password**.

6. Exploring the Gluster Management Console User Interface

This section describes the main sections of the user interface and explains the various functions you can perform using each section.

Note: You can use the Management Console to set up different cloud environments; but can only manage one cloud in one login session.

The Gluster Management Console window is divided into two panes: a navigation tree in the left and a tabbed information pane in the right.



- **Left Pane:** Displays the cloud, registered servers, and volumes. It also displays the discovered servers as well.
- **Right Pane:** At any point of time, the **Right pane** shows tabs with information that are applicable to the resource selected in the left pane such as Server, Volumes, Bricks, and Discovered Server.
- **File Menu:** Helps you to exit the application.
- **Edit Menu:** Helps you to import keys, export keys, change password, change settings and refresh the application.
- **Gluster Menu:** Helps you to manage your cluster, servers, and volumes. For example, if you have selected volumes, the Gluster menu displays all volume related options.

6.1. Gluster Management Console Toolbar

The Toolbar provides shortcuts to the main Gluster Management Console functions.

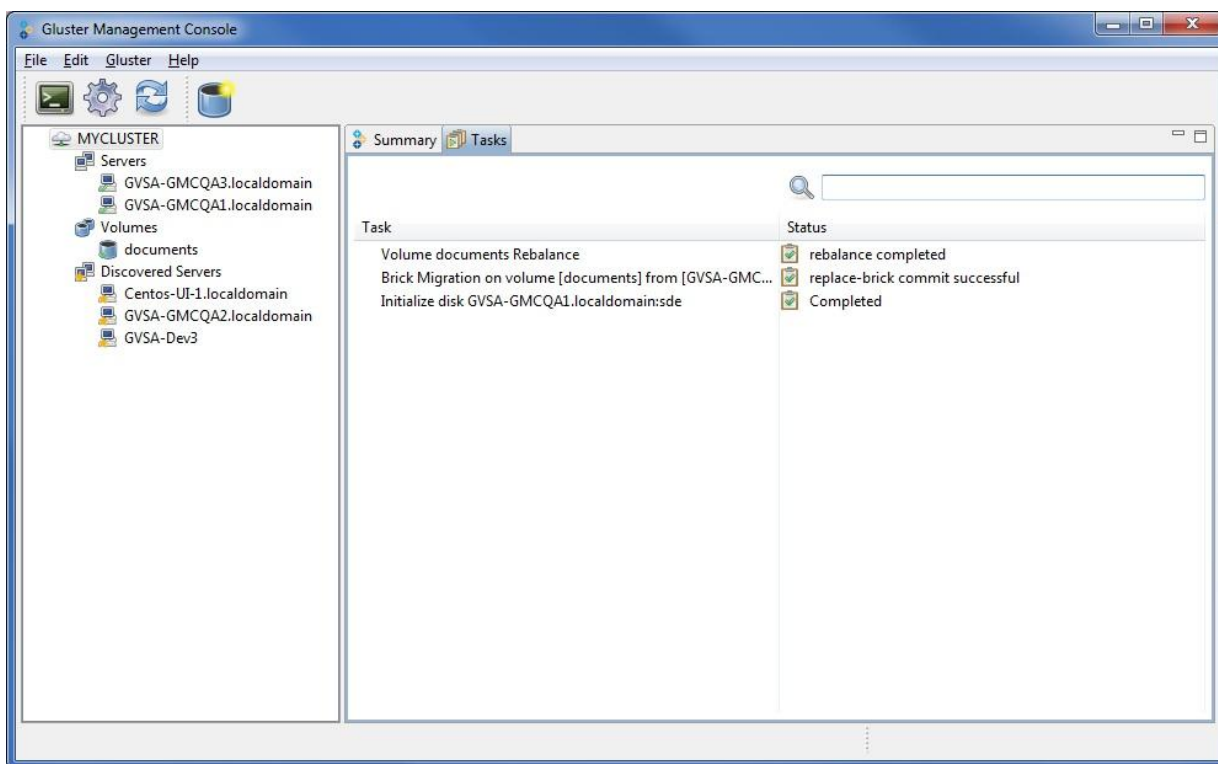
The table below describes the buttons on the toolbar of the console.

| Icon | Description |
|---|--|
|  | Launches the terminal window that allows you type system commands. |
|  | Opens the Preferences window. |
|  | Refreshes the application. |
|  | Allows you to add a server. |
|  | Allows you to remove a server. |
|  | Opens the Create Volume window. |
|  | Allows you to start a volume. |
|  | Allows you to stop a volume. |
|  | Allows you to remove a volume. |
|  | Allows you to rebalance a volume. |
|  | Allows you to reset the volume options. |
|  | Opens the Add Brick window. |
|  | Removes brick from a volume. |
|  | Migrates bricks from one volume to another |
|  | Allows you to download logs. |
|  | Allows you to resume a task from the Tasks tab. |

| Icon | Description |
|------|--|
| | Allows you to stop a task from the Tasks tab. |
| | Allows you to pause a task from the Tasks tab. |
| | Allows you to clear completed task from the Tasks tab. |

6.2. Tasks

Task tab is displayed when you select the cluster node from the navigation tree. It displays the status of all running tasks such as disk initialization, brick migration, volume rebalance.



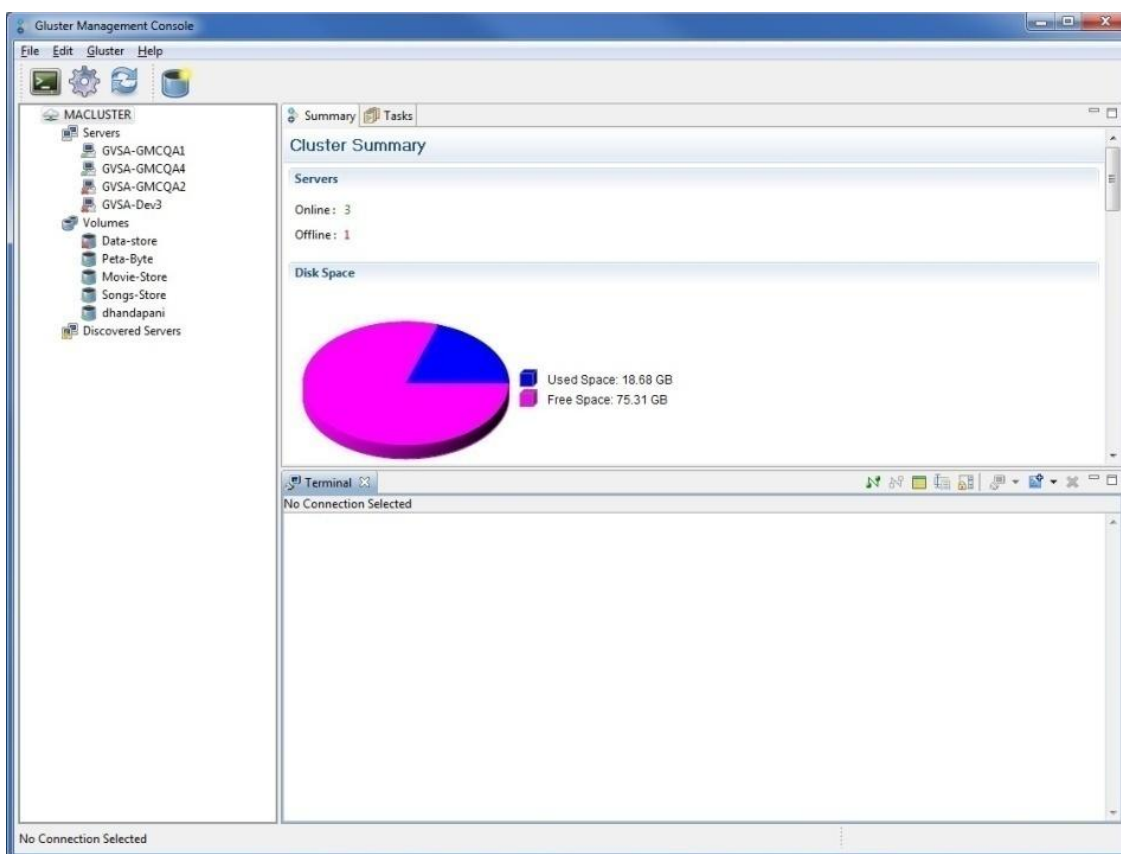
When you select a task on this tab, the relevant toolbar options are enabled such as pause, resume, abort, commit, or clear.

6.3. Terminal


The Gluster Management Console provides a stand-alone ANSI or vt102 terminal emulator that enables you to connect to servers through SSH.

It processes ANSI control characters, including NUL, backspace, carriage return, linefeed, and a subset of ANSI escape sequences sufficient to enable the use of screen-oriented applications, such as vi, Emacs, and any GNU readline-enabled application such as Bash, bc, or ncftp.

Note: The Terminal window is not a fully compliant vt100 or vt102 terminal emulator.



You must provide SSH connection details to connect to the required host. Here is a screenshot of how to configure these settings on Gluster Management Console:



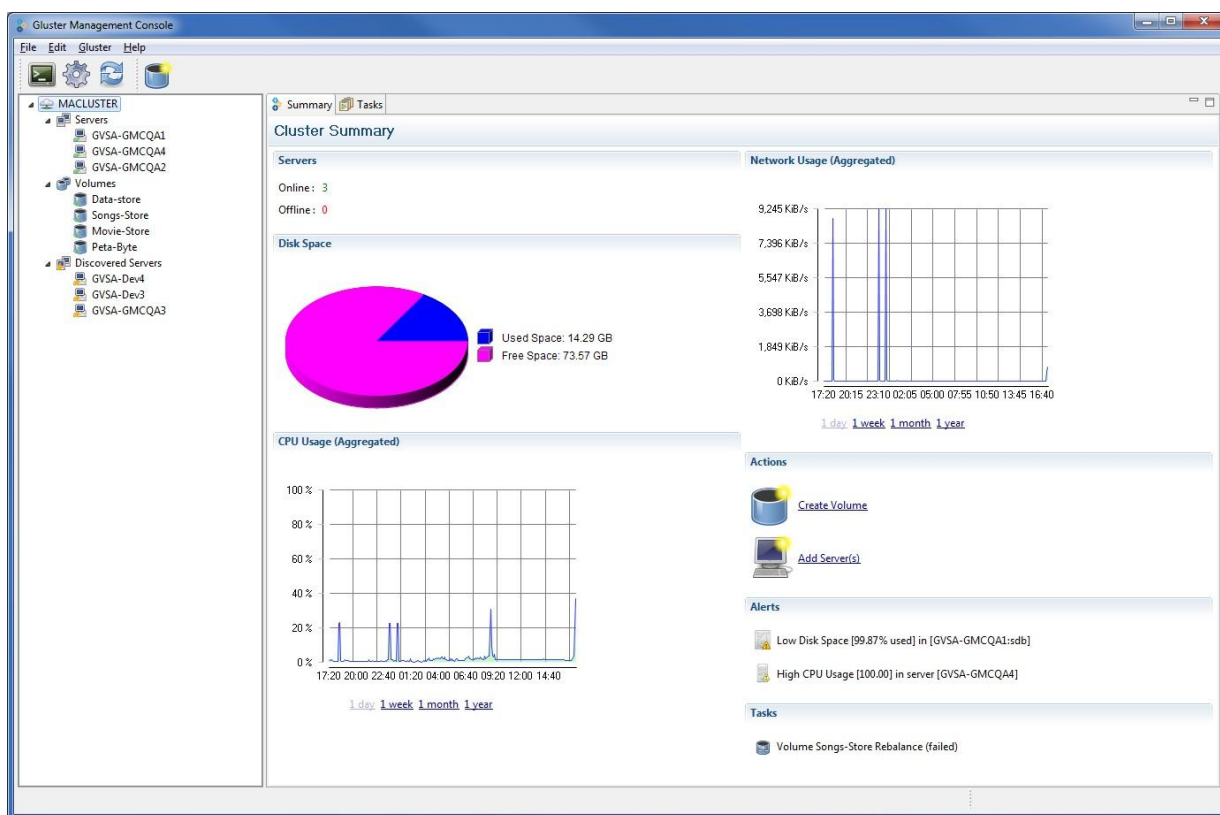
6.4. Monitoring Cluster Resources

Gluster Management Console enables you to monitor your cluster resources such as servers, volumes, through **Summary** tabs.

To monitor resources in your cluster, perform the following steps.

1. Select the resource you wish to monitor from the left pane.

For example, choose your cluster. The **Summary** tab displays the historical graphical representations of your cluster's key resources such as CPU usage, memory, network usage, and disk utilization by day, week, or month. In case of Network usage, the graph will be displayed for the selected network interface from the drop down list.



2. Click to expand other resource nodes such as Volume, Bricks, and Server to display their respective details, as required.

6.5. Alerts

Alerts section lists the typical alarm conditions found in your cluster environment. For example, low disk space is frequently a problem on managed machines. You can generate alerts for the critical conditions in your cluster by defining threshold conditions. An alert is generated when it reaches the threshold conditions that you set against server resources. These alerts are real-time data collected from your resources and are displayed in the **Cluster Summary** tab.

- CPU utilization
- Memory utilization
- Disk-space utilization
- Server offline
- Brick offline

To know more on how you can set the thresholds conditions, refer to [Section 5.3.2. Setting Alerts Preferences](#).

6.6. Preferences


You can use preferences to configure several aspects of the Gluster Management Console. These include setting preferences for the console, alerts, charts, and terminal.

6.6.1. Setting Gluster Preferences

To set the Gluster preferences, perform the following steps:

1. Choose **Edit > Settings > Gluster**

The **Preferences** window displays the **Gluster** page on the right pane of the window.



2. Select **Show Cluster Selection Dialog on Login**, if you want your cluster selection dialog to appear when you launch the application.
3. Choose your cluster name from **Default Cluster to Manage** drop down list to manage the default cluster you wish to manage when you launch the application.
4. Type the duration (in sec) of the refresh intervals in the **Refresh Interval** field.
5. Click **Apply**.
6. Click **OK** to exit the **Preferences** window.

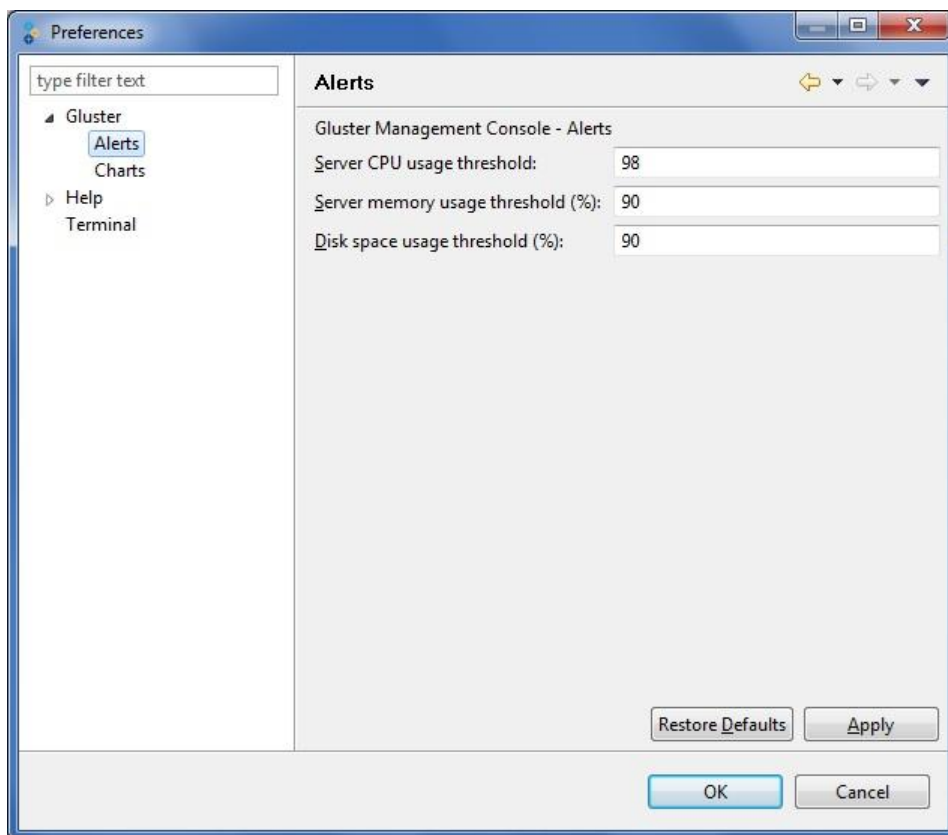
6.6.2. Setting Alerts Preferences

Thresholds are the critical conditions that you define in Gluster Management Console.

To set the thresholds, perform the following steps:

1. Choose **Edit > Settings > Alerts**

The **Preferences** window displays the **Alerts** page on the right pane of the window.



2. Type the value for server CPU threshold limits in the **Server CPU Threshold** field.
3. Type the value for server memory usage threshold in the **Server Memory Usage Threshold** field.
4. Type the value for disk space usage threshold in the **Disk Space Threshold** field.

5. Click **Apply**.
6. Click **OK** to exit the **Preferences** window.

6.6.3. *Setting Charts Preferences*

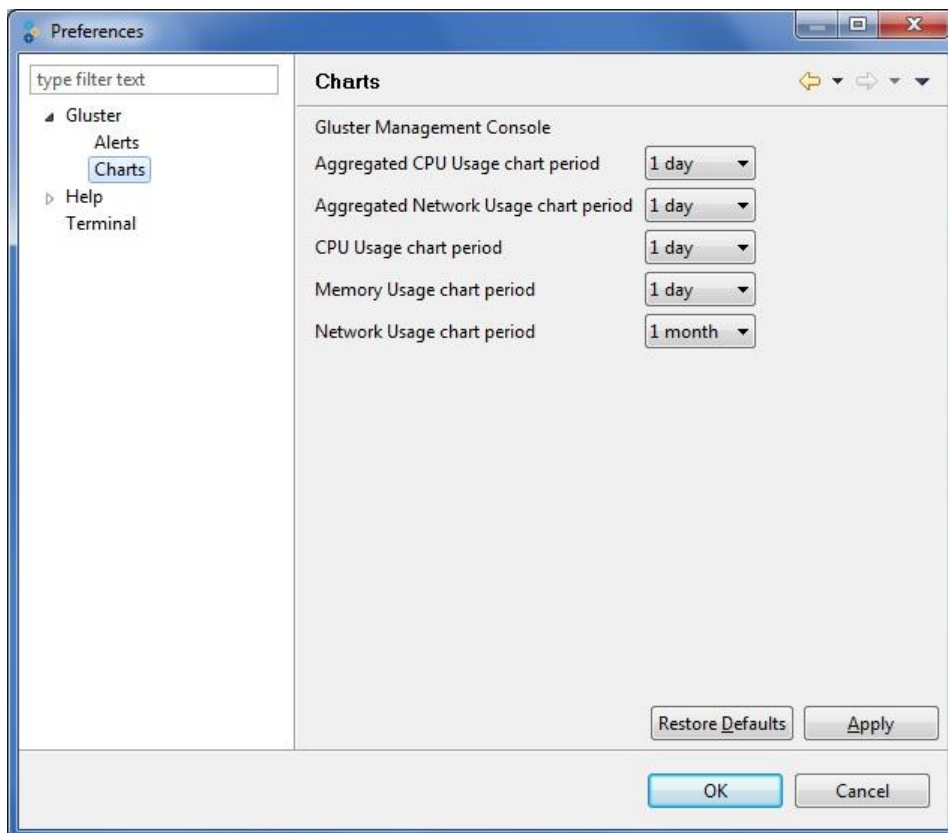
You can customize the console and change the way some of the monitoring data is displayed, as described in this section.

Note: The default duration for all the charts is a day.

To set the display of the charts, perform the following:

1. Choose **Edit > Settings > Charts**

The **Preferences** window displays the **Charts** page on the right pane of the window.



2. Select the **Aggregated CPU Usage chart period** from the drop down list.
3. Select the **Aggregated Network Usage chart period** from the drop down list.
4. Select the **CPU Usage chart period** from the drop down list.
5. Select the **Memory Usage chart period** from the drop down list.
6. Select the **Network Usage chart period** from the drop down list.

7. Click **Apply**.
8. Click **OK** to exit the **Preferences** window.

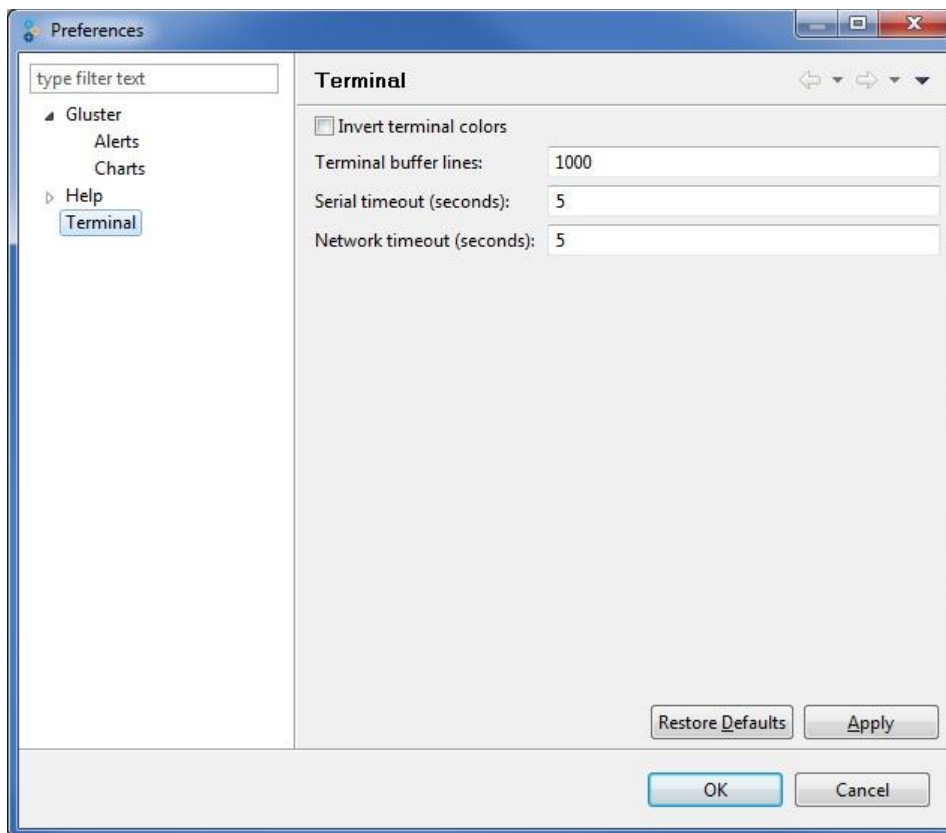
6.6.4. *Setting Terminal Preferences*

You can customize the terminal window as described in this section.

To set the display of the charts, perform the following:

1. Choose **Edit > Settings > Terminal**

The **Preferences** window displays the **Terminal** page on the right pane of the window.



2. Select **Invert terminal colors** check box if you wish to invert the color of the terminal.
3. Type the limit of terminal buffer lines in the **Terminal buffer lines** field.
4. Type the value of the serial timeout in the **Serial timeout** field.
5. Type the value of the network timeout in the **Network timeout** field.
6. Click **Apply**.
7. Click **OK** to exit the **Preferences** window.

6.7. Security Keys

Gluster Management Gateway communicates with storage servers using public key authentication. This establishes trust between your account on your central server and the new machine you just added to your cluster. Administrators can export and/or import the SSH keys (of the gateway) for security reasons from the Gluster Management Gateway server.

6.7.1. Importing Keys

To import the security keys, perform the following steps:

1. Choose **Edit > Import Keys**.
2. A window is displayed allowing you to choose the download location on your machine.
3. Choose the location of the `.tar` file containing the keys, and click **Save**.

The security keys are imported from the gateway server to the location you specified in `ssh-keys.tar` file format.

6.7.2. Exporting Keys

To export the security keys, perform the following steps:

1. Choose **Edit > Export Keys**.
2. A window is displayed allowing you to choose the location on your machine.
3. Choose the location, and click **Save**.

The security keys are exported from the gateway server to the location you specified in a `.tar` file format.

6.8. Resetting the Password

If you have managed to get locked out of your Gluster Management Console, you will be able to reset the password if you are the administrator of the Tomcat Server.

Note: Ensure to stop the Tomcat5 service before running this script.

You can reset the password by running the following tool on your Tomcat server:

```
gmg-reset-password.sh
```

The tool resets the password to the default system password "`sys3m`".

6.9. CIFS Services

Common Internet File System (CIFS) operates as an application-layer network protocol to provide shared access to files, printers, serial ports, and miscellaneous communications between nodes on a network. It also provides an authenticated inter-process communication mechanism. Most usage of

SMB involves computers running Microsoft Windows operating system.

CIFS can be used as an access protocol by Microsoft Windows users to access files in a GlusterFS volume. Users will be able to perform regular file system operations such as creating files, read or write data to files, deletion of data through CIFS.

Typically, all CIFS related operations such as configuring CIFS services, creating and removing users, enabling CIFS re-export of a volume must be performed on each and every node of your cluster. Gluster Management Gateway provides an easy way to manage these operations, such as:

- An user interface for enabling or disabling CIFS re-export of volumes and configuring CIFS users for a volume
- Command line utilities for configuring CIFS services, creating and removing CIFS users on all the servers.

6.9.1. Configuring the CIFS Services

To configure CIFS, you must create a text file consisting of a list of servers (hostnames or IP addresses) on which CIFS needs to be configured. Ensure that the text file contains one server name per line.

Run the following command on the gateway server to configure CIFS:

```
$ setup_cifs_config_all.py <server list file>
```

Note: Ensure that CIFS services are configured on **all** the servers of your cluster at all times. Any new servers added to the cluster, ensure that you configure CIFS services on them. This enables you to access the CIFS enabled volumes from any of the cluster nodes.

6.9.2. Creating CIFS Users

You must create users on each of the servers to enable them to use the CIFS access protocol.

To create CIFS users, run the following command on the gateway server:

```
$ add_user_cifs_all.py <server list file> <user name> <password>
```

- Server list file - contains one server name per line

The user defined in the *<user name>* will be created on each and every server defined in *<server list file>*, will be assigned the same password defined in *<password>*.

Note: Ensure that the CIFS users are created on **all** the servers of your cluster at all times. Any new servers added to the cluster, ensure that you create the CIFS users on them.

6.9.3. Removing CIFS Users

To delete CIFS users, run the following on the gateway server:

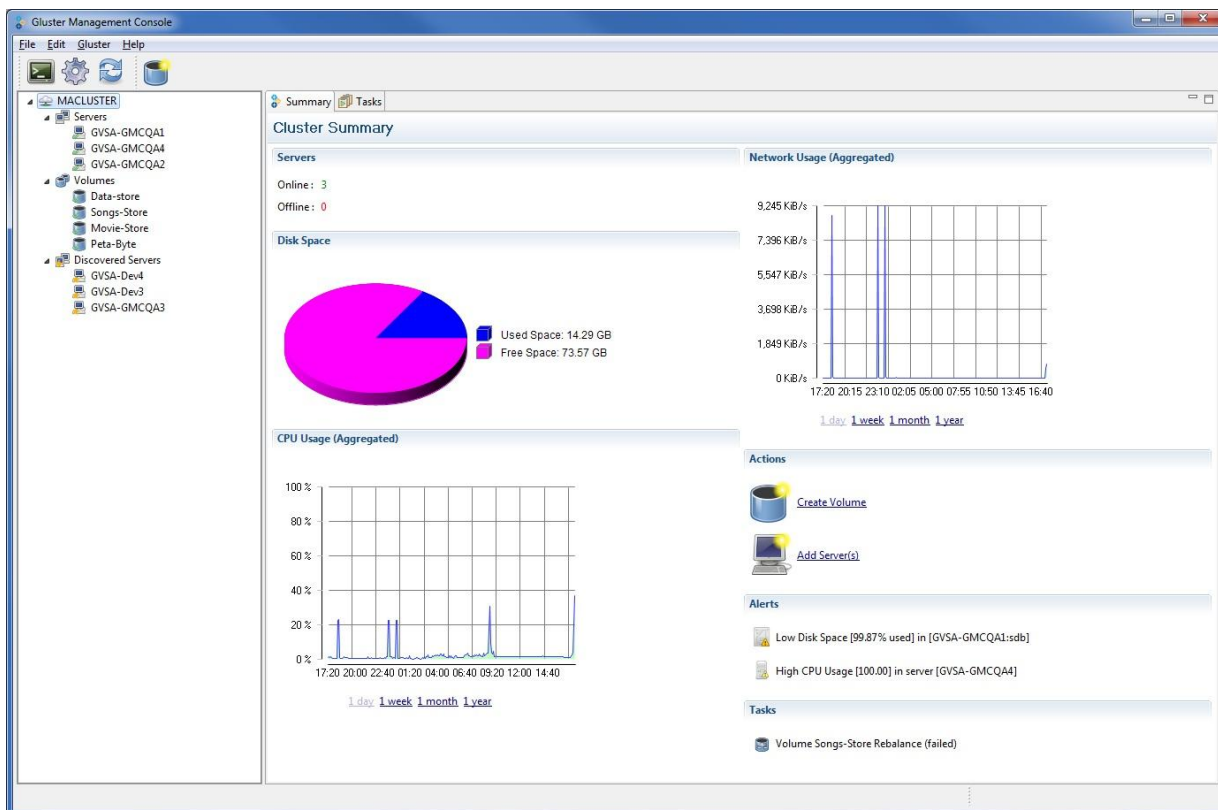
```
$ delete_user_cifs_all.py <server list file> <user name>
```

7. Managing Clusters

A cluster is a group of linked computers, working together closely thus in many respects forming a single computer. Gluster Management Console enables you to create and manage your clusters.

Cluster Summary

The console enables you to monitor your cluster resources from the **Summary** tab. To display the summary of your cluster, click the cluster node from the left pane of the console window. The summary of the selected cluster is displayed in the right pane of the console window.



After you log into the application successfully, you will be prompted to:

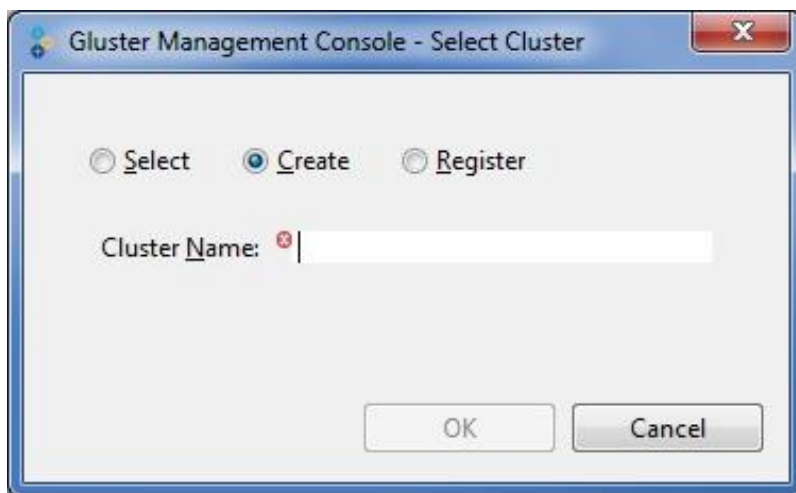
- Select Cluster - The cluster you will manage in the login session
- Create Cluster - Creating an empty cluster that can be expanded using the console.
- Register Cluster - Registering an existing cluster and start managing it through the console.

7.1. Creating a Cluster

To create a cluster, perform the following steps.

1. Launch the application and log in with a valid password.

Select Cluster window is displayed.



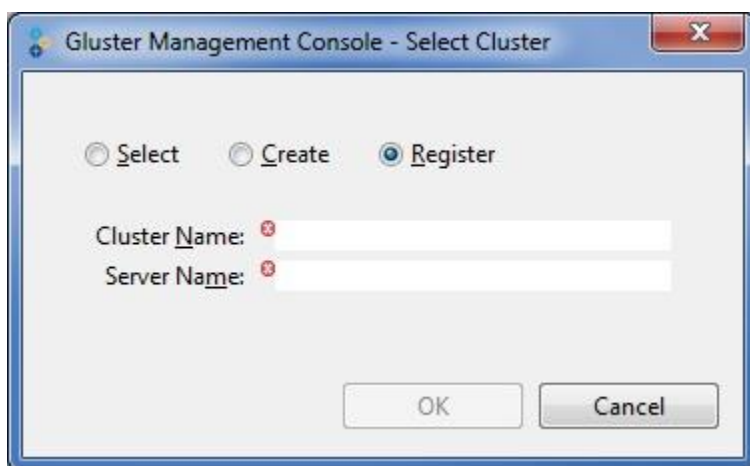
2. Select **Create**. Type the name of the new cluster in the **Cluster Name** field.
3. Click **OK**. The **Gluster Management Console** window is displayed with the newly created cluster.

7.2. Registering a Cluster

To register a cluster, perform the following steps.

1. Launch the application and log in with a valid password.

Select Cluster window is displayed.



2. Select **Register**. Type the name of the cluster and server name in the respective fields.

Note: You must type the name of the server which is online in the **Server Name** field.

3. Click OK. The Gluster Management Console will register the cluster with the gateway; fetch all the management information such as servers, volume resources and display them in the console.


7.3. Selecting a Cluster

The select cluster option is enabled after you create or register a cluster.


To select a cluster, perform the following steps.

1. Launch the application and log in with a valid password.

Select Cluster window is displayed.



2. Choose the cluster name from the Cluster Name drop-down list.
3. Click OK. **Progress Information** window is displayed while the console fetches the management information for the selected cluster.



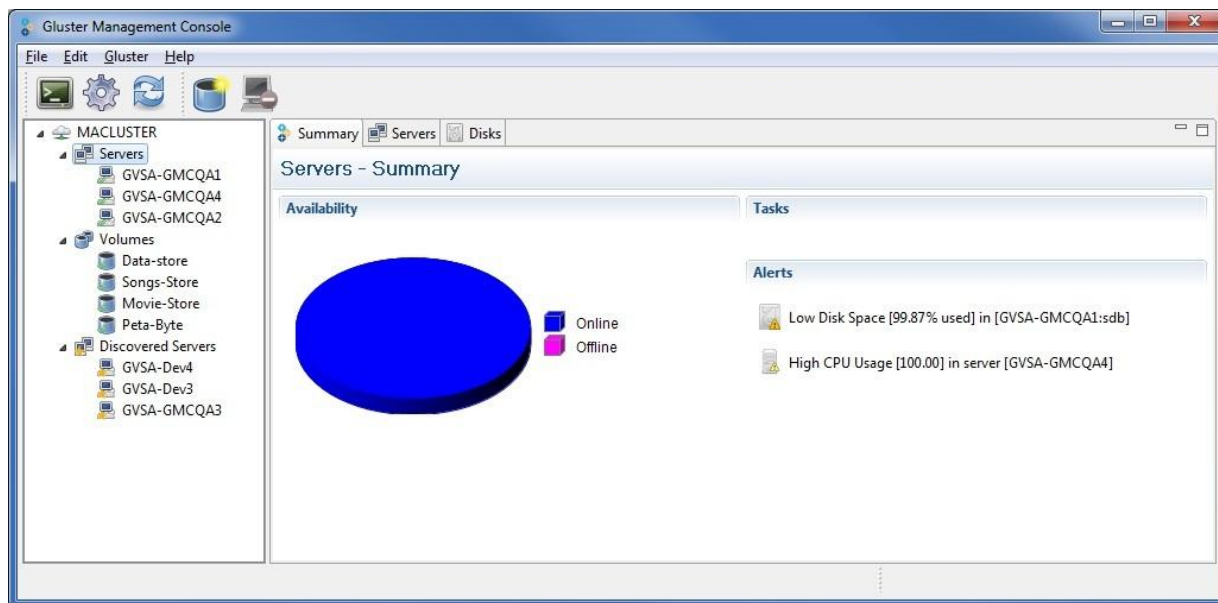
4. The **Gluster Management Console** window is displayed with the selected cluster.

8. Managing Servers

A server is a computer on a network that is used as a storage node in a cluster. You can add a new server to the storage cluster to expand the amount of available storage.

Server Summary

The console enables you to monitor your servers in your cluster from the **Summary** tab. To display the summary of your servers, click the **Servers** node from the left pane of the console window. **Servers** tab provides a summary of all the servers, list of running tasks, and any alerts of your servers in your cluster. The **Disks** tab provides a summary of all the disks on all the servers of your cluster.



This section describes how you can manage your servers in your storage cluster using the management console.

- [Adding Servers to a Cluster](#)
- [Removing Servers from a Cluster](#)
- [Initializing a Disk](#)

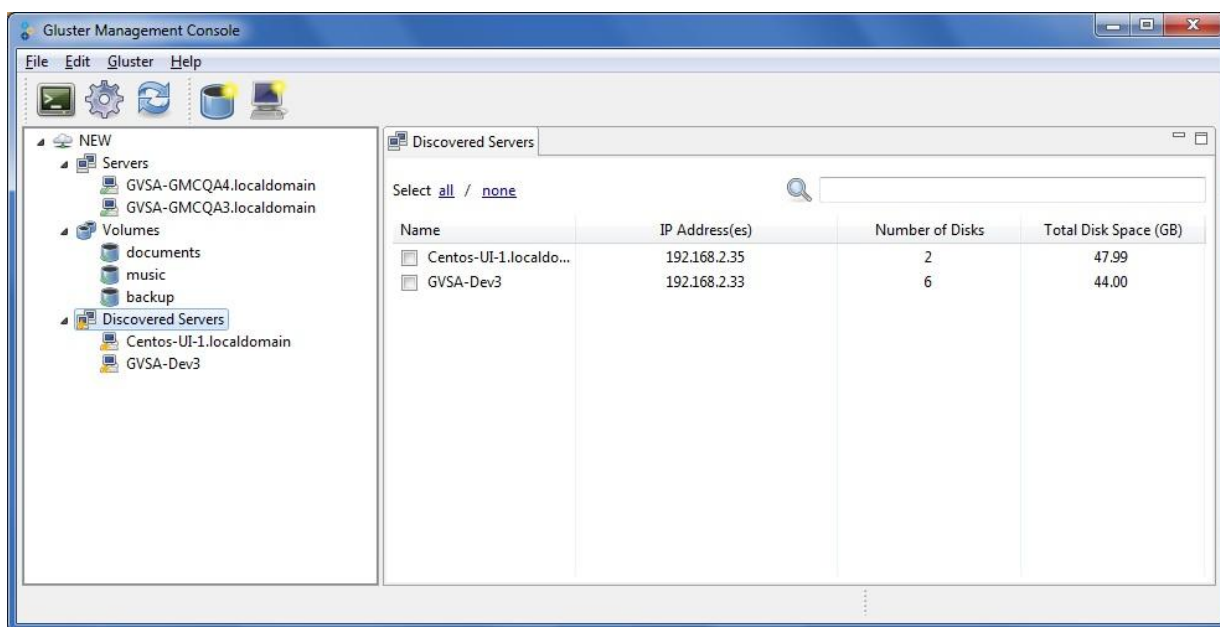
8.1. Adding Servers to a Cluster

Gluster Management Console discovers servers across the environment. You can add servers manually or from the discovered list of servers to your cloud environment to expand the amount of available storage.

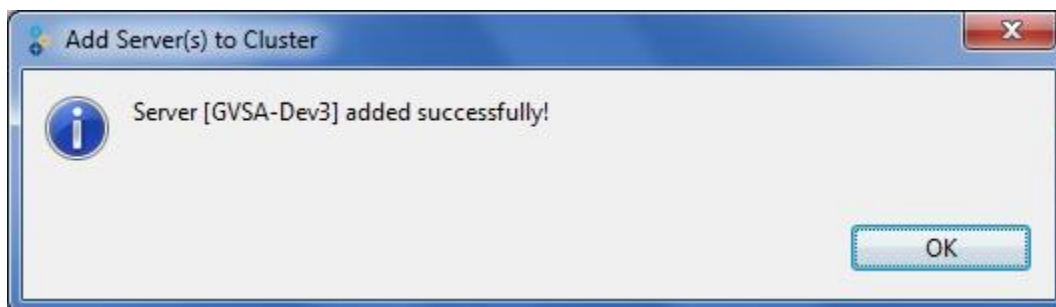
To add a discovered server, perform the following steps:

1. Choose from the menu, toolbar, or right click on the Discovered Server node and select **Add Server**.

You can also choose multiple discovered servers or select **All** from the **Discovered Servers** tab and select **Add Server**.



2. The console displays the **Add Server(s) to Cluster** window to indicate that the server was successfully added to the cluster.



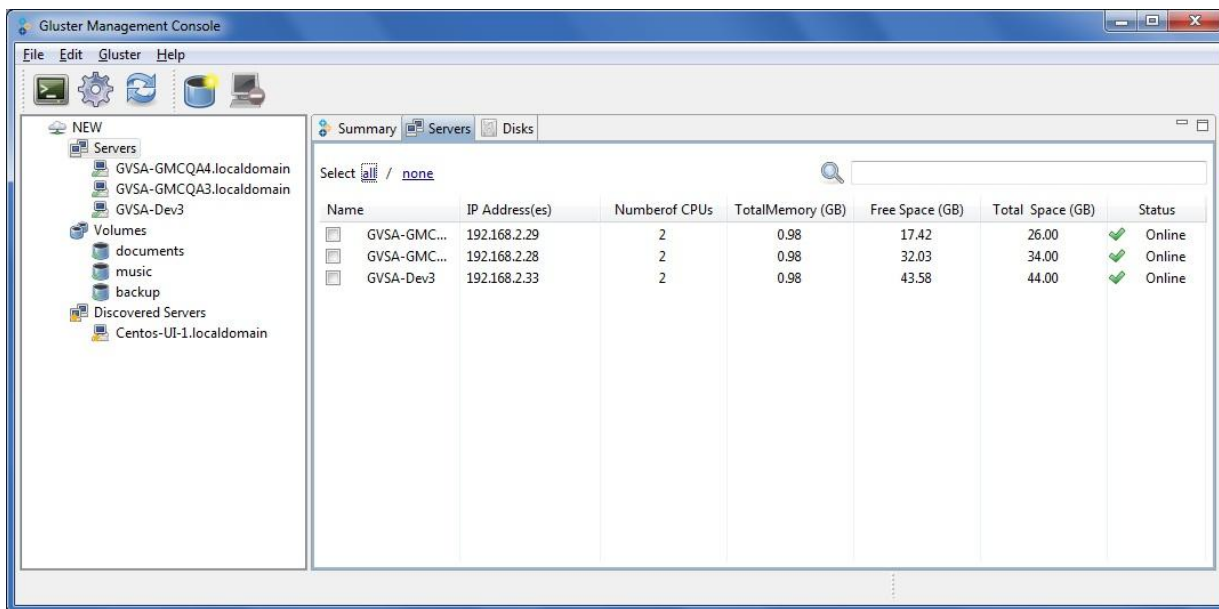
3. Click **OK**.

8.2. Removing Servers from a Cluster


To remove servers from a cluster, perform the following steps:

1. Choose from the menu, toolbar, or right click on the selected server. Select the name of the server you want to remove from a cluster and click **Remove Server**.


You can also choose multiple servers or select **All** from the **Servers** tab and select **Remove Server**.



You are asked to confirm before removing a server.



2. Click **Yes**. The console displays the following window to indicate that the server was successfully removed.



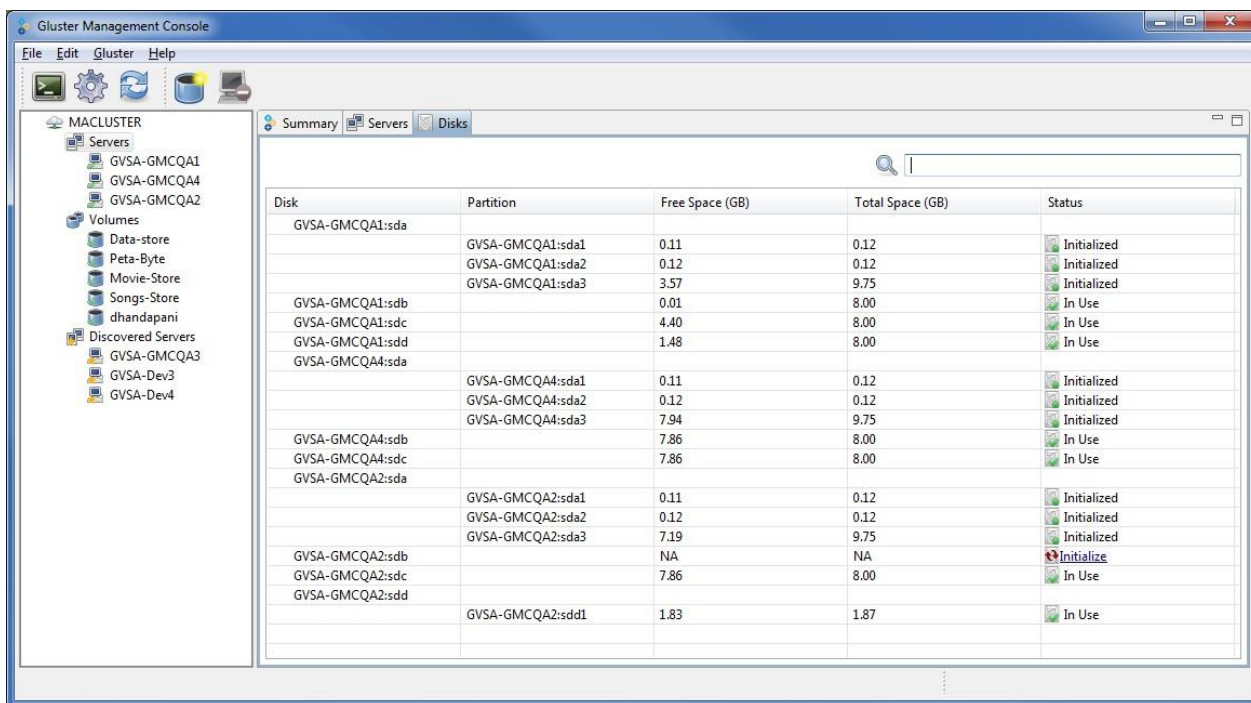
3. Click **OK**. The removed server is listed under **Discovered Server** node.

8.3. Initializing a Disk

To initialize a disk, perform the following steps.

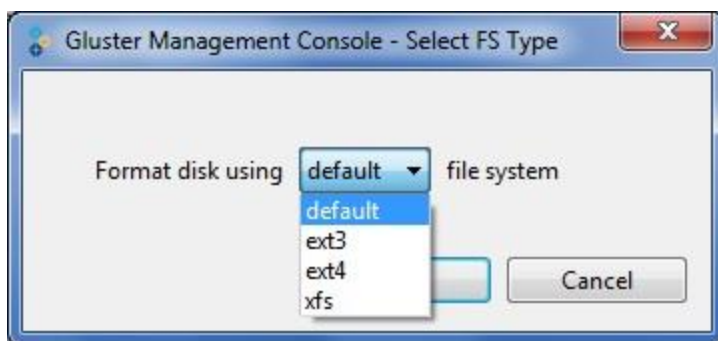
1. In the left pane of the console window, select the **Server** or **Servers** node.

The **Disks** tab displays all the disks in your cluster or selected server. The tab displays the **Initialize** link for all uninitialized disks in the **Status** column.



2. Select the disk to format and click **Initialize**.

The **Select Filesystem Type** window is displayed.



3. Choose the **File System** type from the drop-down list. Click **OK**.

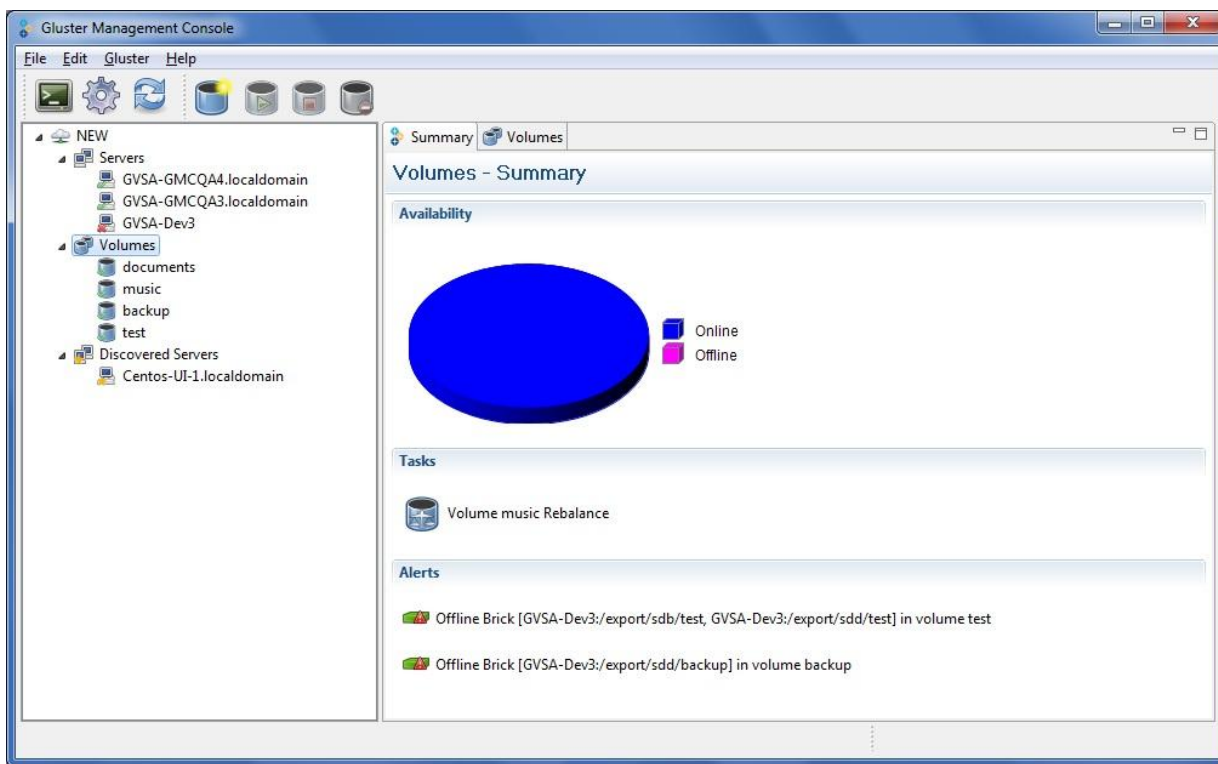
The selected disk begins formatting and displays the current status in the **Task** tab.

9. 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 server in the trusted storage pool. Most of the Gluster management operations happen on the volume. The volume is the designated unit of administration in GlusterFS, so managing them is a large part of the administrator's duties.

Volume Summary

The console also enables you to monitor the volumes in your cluster from the **Summary** tab. To display the summary of your volumes, click the Volume node from the left pane of the console window. The summary of the volumes is displayed in the right pane of the console window. It also displays **Tasks** and **Alerts** of all the volumes.



This section explains how to manage the volumes stored on the file server machines.

- [Creating a Volume](#)
- [Starting Volumes](#)
- [Tuning Volume Options](#)
- [Stopping Volumes](#)
- [Rebalancing Volume](#)
- [Deleting Volumes](#)

- [Managing Bricks](#)
 - [Adding Bricks](#)
 - [Removing Bricks](#)
 - [Migrating Bricks](#)

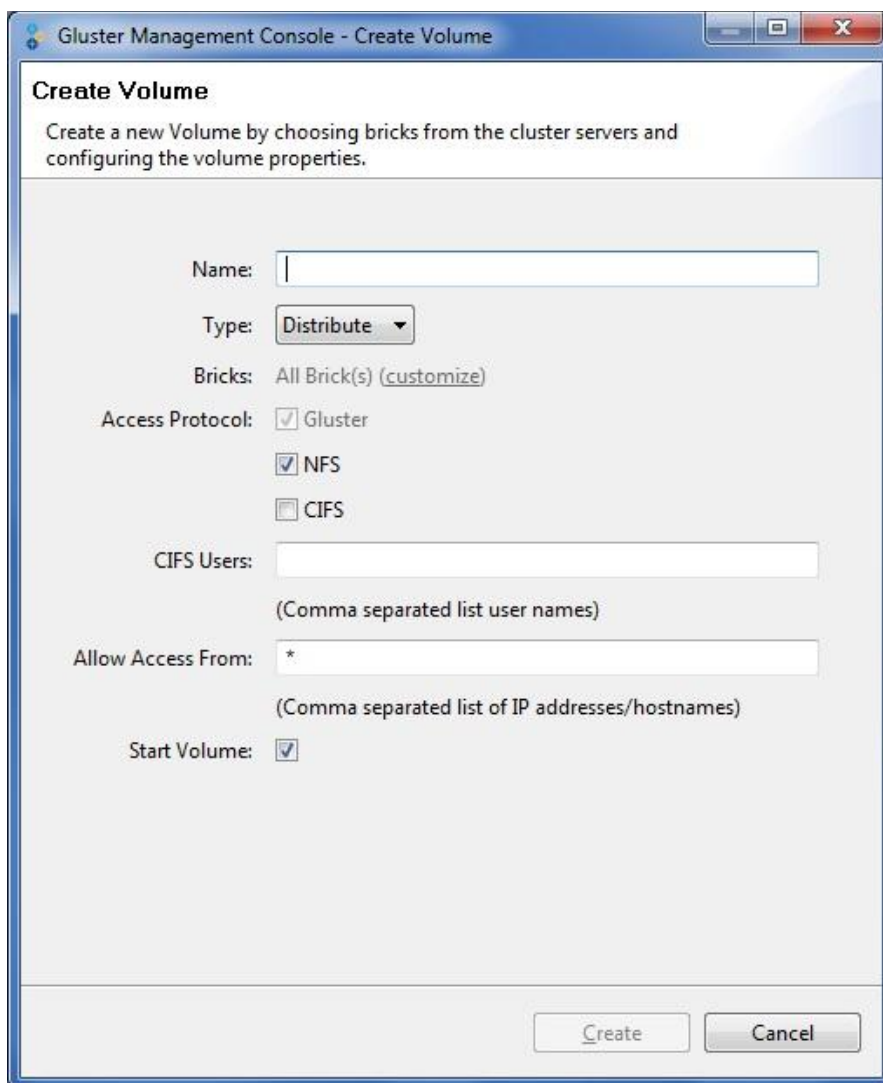
9.1. Creating a Volume

You can create new volumes in your storage environment, as needed. When creating a new volume, you must specify the bricks that comprise the volume and specify whether the volume is to be distributed, replicated, or striped.

To create a volume, perform the following steps:

1. Choose from the menu, toolbar, or right click on the Server and select **Create Volume**.

The **Create Volume** window is displayed.



2. Type the name of the volume in the **Name** field.
3. Select the type of the volume from the **Type** drop-down list. You can choose the volume type as Distribute, Replicate, or Stripe.
4. Bricks will be created on all the disks of the cluster when you create a volume, by default. Click **Customize**, to select and add specific bricks to your volume.
5. Select the **Access Protocol** for the new volume by selecting **GlusterFS**, **NFS**, or **CIFS** checkbox.
6. If you choose **CIFS**, specify the user name and password to use from Windows clients to connect to the volume.

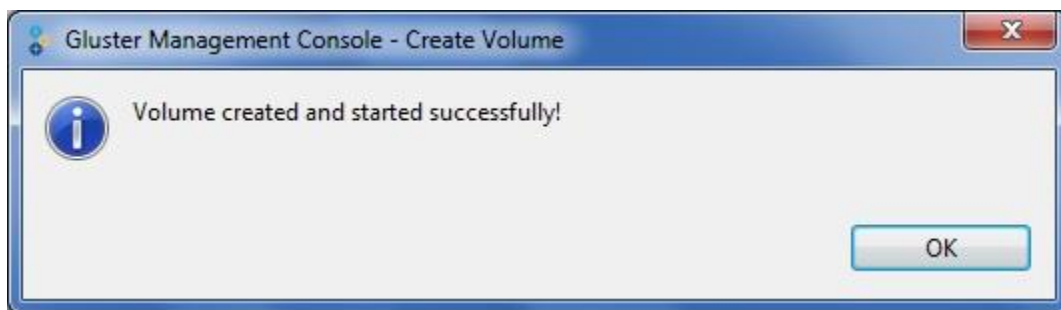
You need to use user-based authentication for CIFS exports, hence you must specify the users who can access the volume as a comma-separated list in the **CIFS Users** field.

For more information about configuring CIFS, refer to [Section 5.8, CIFS Services](#)

7. 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. Typing an asterisk (*) specifies all IP addresses or hostnames. You need to use IP-based authentication for Gluster Filesystem and NFS exports.

8. Select the **Start Volume** checkbox to automatically start the volume once created.
9. Click **Create Volume**. The console displays **Create Volume** window to indicate that the volume was successfully created.



10. Click **OK**.

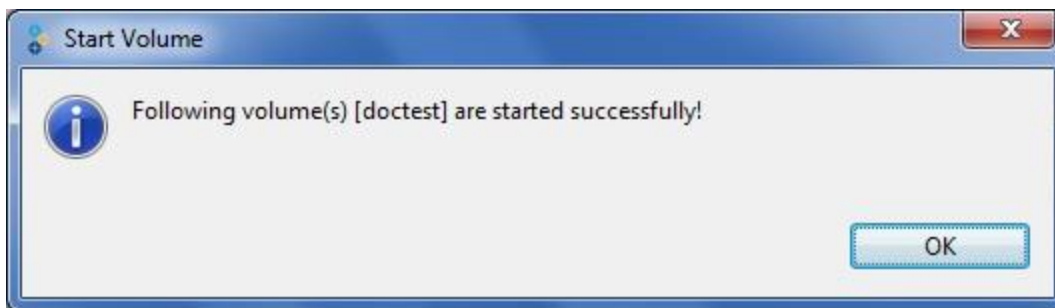
9.2. Starting Volumes

You can start a volume or multiple volumes that are currently stopped. To start volumes, perform the following steps:


1. Choose from the menu, toolbar, or right click to select the volume and select **Start Volume**.

Alternatively, from the left pane you can choose the **Volumes** node and select the volumes that you need to start in **Volumes** tab or click **All** to select all volumes and select **Start Volume**.

2. The **Start Volume** window appears indicating that the volume was successfully started.



If you have selected multiple volumes, the Start Volume window appears indicating the volumes were started successfully.

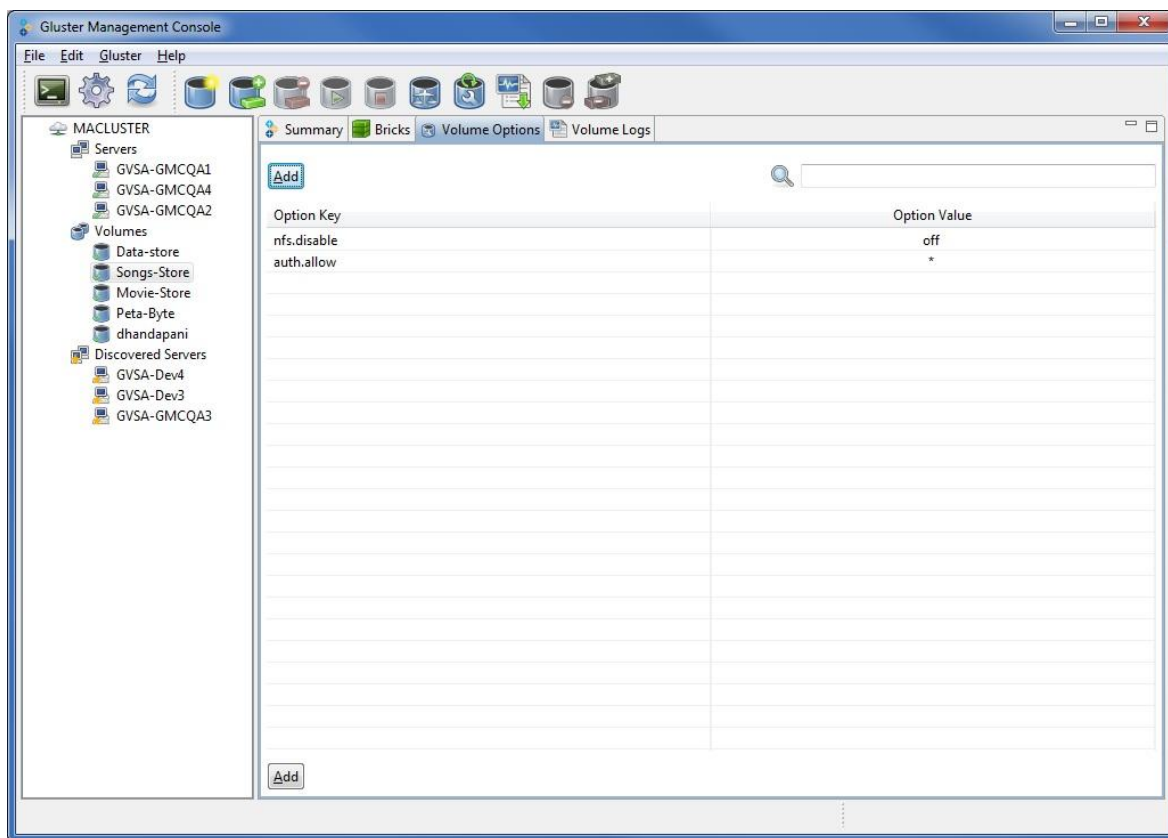


3. Click **OK**.

9.3. Tuning Volume Options

To tune volume options, perform the following steps.

1. Select the volume for which you need to tune volume options.
2. Click the **Volume Options** tab.



3. Click **Add**. Choose the volume option from **Option Key** drop down list and type the value in the **Option Value** field.

You can edit the value of an existing option by selecting the **Option Value** field.

To read more about the volume options, refer to the [Gluster Filesystem Administration Guide, Section 8.1 Tuning Volume Options](#).

9.4. Stopping Volumes

To stop a volume, perform the following steps:

1. Choose from the menu, toolbar, or right click to select the volume. Select **Stop Volume**.


You are asked to confirm before stopping a volume.



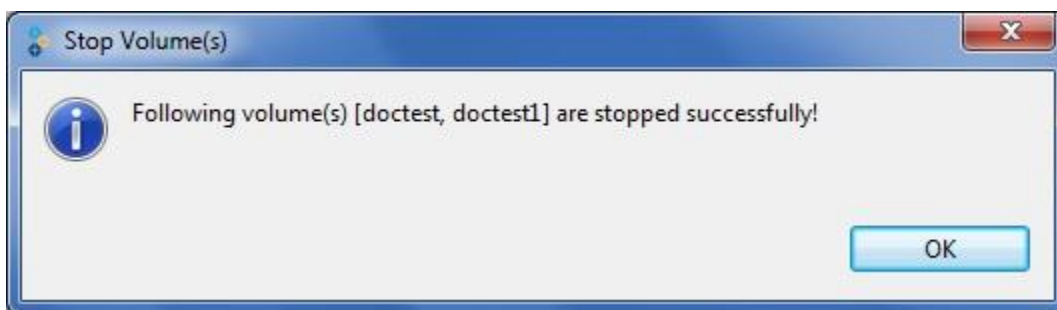
In case of multiple volumes, the **Stop Volume** window is displayed as shown below.



2. Select **Yes** to stop the Volume. The **Stop Volume** window appears indicating that the volume was successfully stopped.



In case of multiple volumes, the **Stop Volume** window is displayed as follows.



3. Click **OK**.

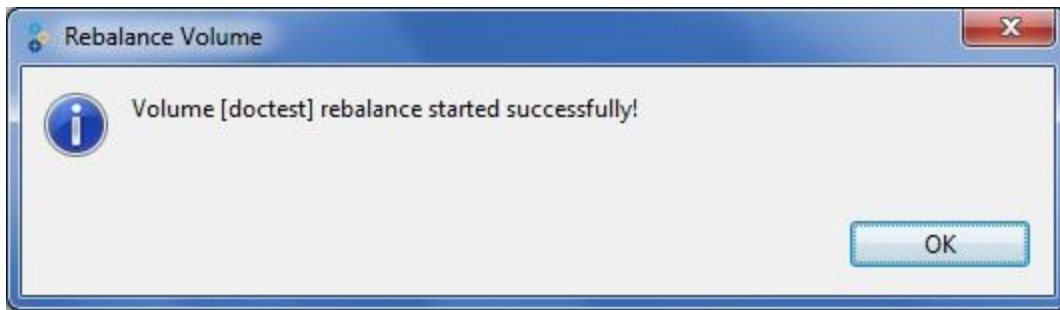
9.5. Rebalancing Volume

You can rebalance a volume, as required. After adding bricks, you need to rebalance the data among the servers. New directories created after expanding the volume will be evenly distributed automatically through this process.

To rebalance a volume, perform the following steps:

1. Choose from the menu, toolbar, or right click and select the Volume and select **Rebalance Volume**.

The **Rebalance Volume** window is displayed indicating that the volume rebalance is successfully started.



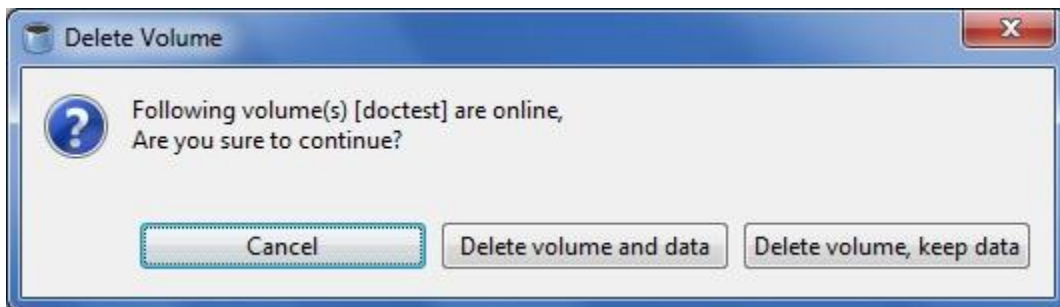
2. Click **OK**.
3. Select the cluster from the left pane and choose the **Task** tab to check the volume rebalance status.
4. If you wish to stop rebalance of a volume, click **Stop** from the menu bar.

9.6. Deleting Volumes

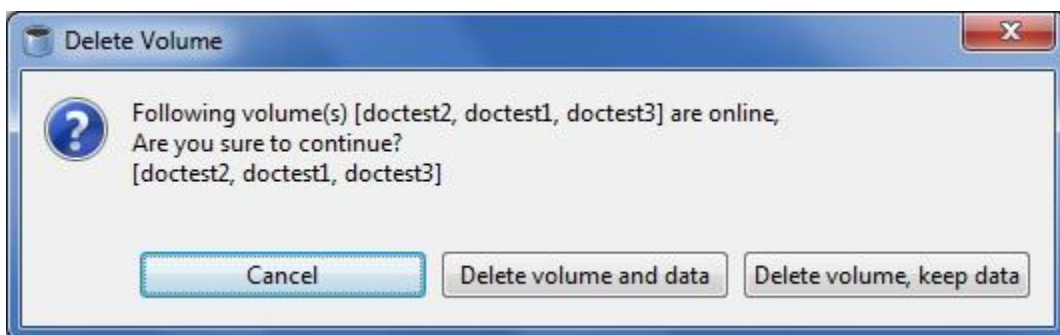
You can delete a volume or multiple volumes from your storage environment, as required. To delete volumes, perform the following steps.

1. Choose from the menu, toolbar, or right click and select the Volume and select **Delete Volume**.

The **Delete Volume window** appears. You are asked to confirm before proceeding with the volume deletion.




In case of multiple volumes, the **Delete Volume** window is displayed as shown below.




2. If you wish to delete the volume permanently, click **Delete Volume and Data**.

- If you wish to keep the data but delete the volume, click **Delete Volume, Keep Data**.
The bricks directories of the volume will be renamed to `<olddirectory>_<timestamp>`
- The **Delete Volume** window is displayed indicating that the volume was deleted successfully.



In case of multiple volumes, the **Delete Volume** window is displayed as follows.



- Click **OK**.

9.7. Managing Bricks

Brick is the basic unit of storage, represented by an export directory on a server in the storage cluster. You can expand, shrink, and migrate volume of your cluster, as required, by adding new bricks, deleting bricks, and migrating bricks.

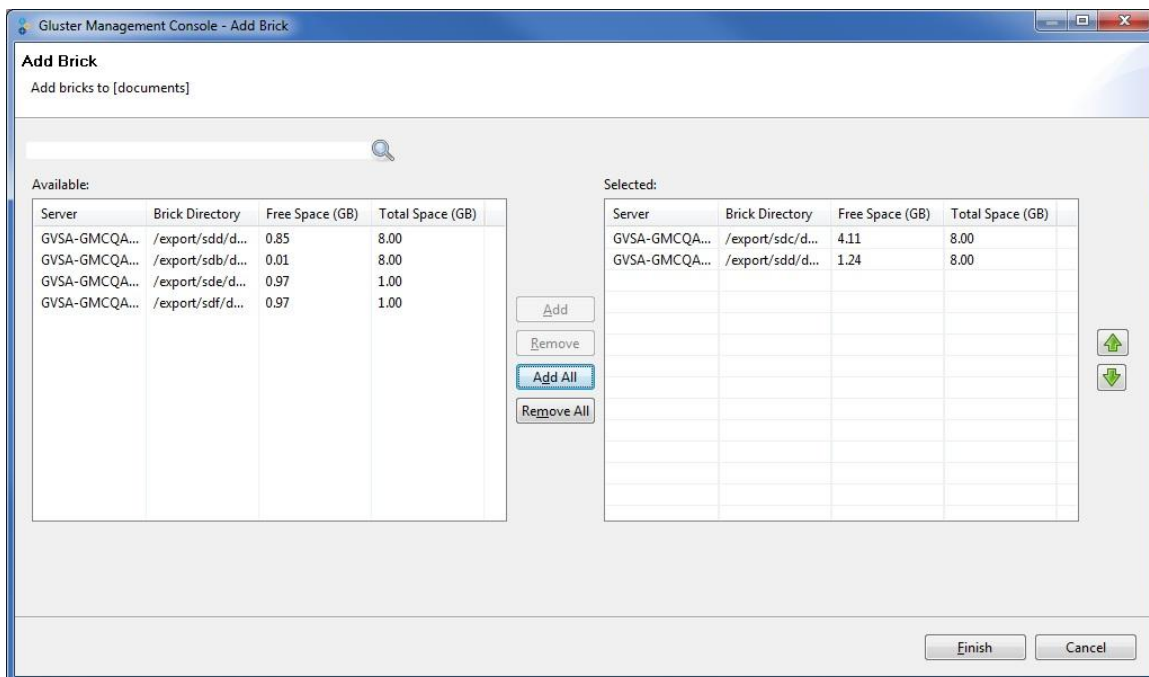
Note: You need to add at least one brick to a distributed volume, multiples of two bricks to replicated volumes, multiples of four bricks to striped volumes when expanding your storage space.

9.7.1. Adding Bricks

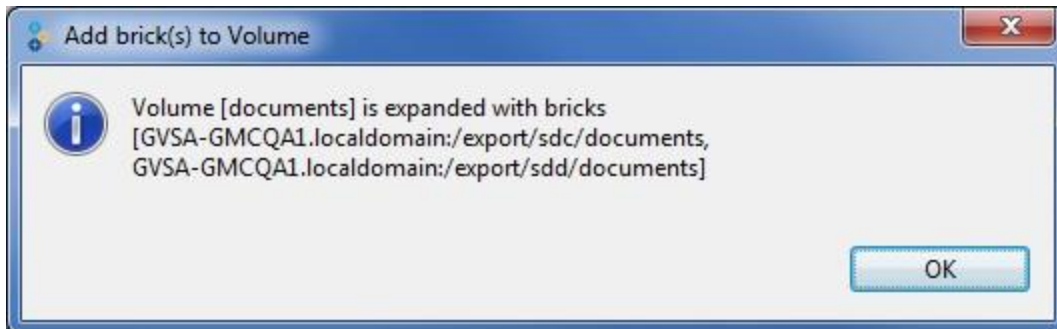
You can expand a volume, as required, by adding new bricks to an existing volume. To add a brick to an existing volume, perform the following steps.

- Select the volume from the left pane, choose from the menu, toolbar, or right click and select **Add Bricks**.

The **Add Brick** window is displayed.



2. Select the bricks that you want to add to the current volume. Click **Add** to add a brick or Click **Add All** to add all the available bricks.
3. Click **Finish**. The **Add Brick(s) to Volume** window appears indicating that the brick(s) is successfully added.



4. Click **Ok**.

9.7.2. Removing Bricks

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

To remove bricks from an existing volume, perform the following steps.

1. Select a brick or bricks from the **Bricks** tab and select **Remove Bricks** from the menu or toolbar.


If the volume is currently running, you are asked to confirm before deleting a Brick.



2. If you wish to delete the brick permanently, click **Remove Bricks, Delete Data**.
3. If you wish to keep the data but delete the brick, click **Remove Bricks, Keep Data**.

The selected bricks directories will be renamed to `<olddirectory>_<timestamp>`.

4. The **Remove Brick** window is displayed indicating that the volume was deleted successfully.



5. Click **Ok**.

9.7.3. Migrating Bricks

You can migrate the data from one brick to another, as needed, while the cluster is online and available.

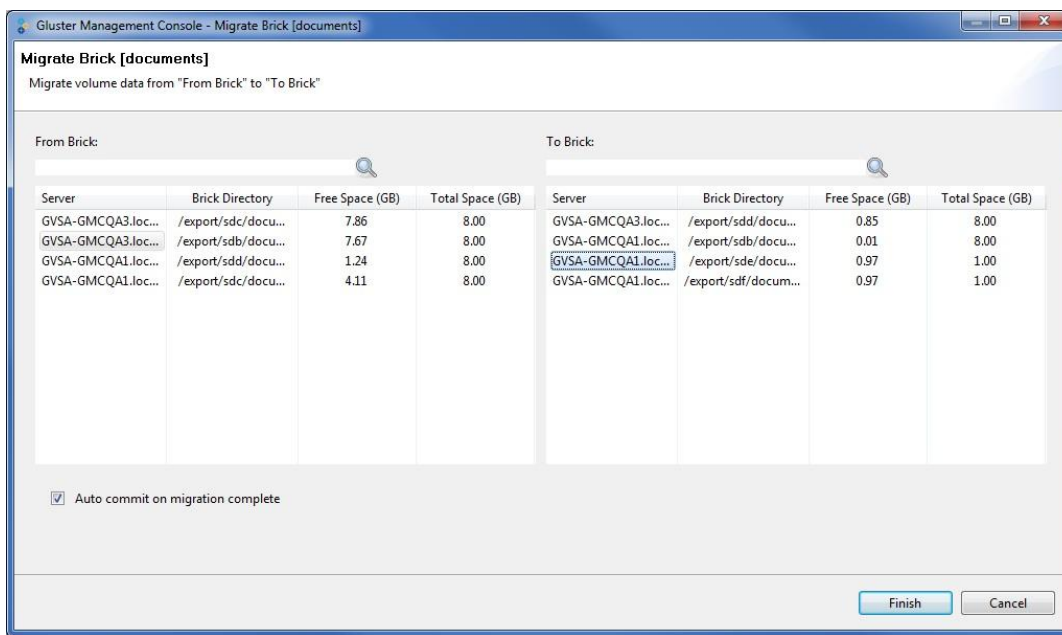
You can migrate bricks, that is, move data from one brick to another within your cluster. You might want to do this, for example, when upgrading a disk to another with a larger capacity.

When migrating a brick, the Gluster Management Console allows you to select both the source and destination bricks. After initiating a disk migration, the **Task** tab displays the status of the migration. You can stop, pause, resume, or commit the migration by selecting the appropriate icon from the toolbar.

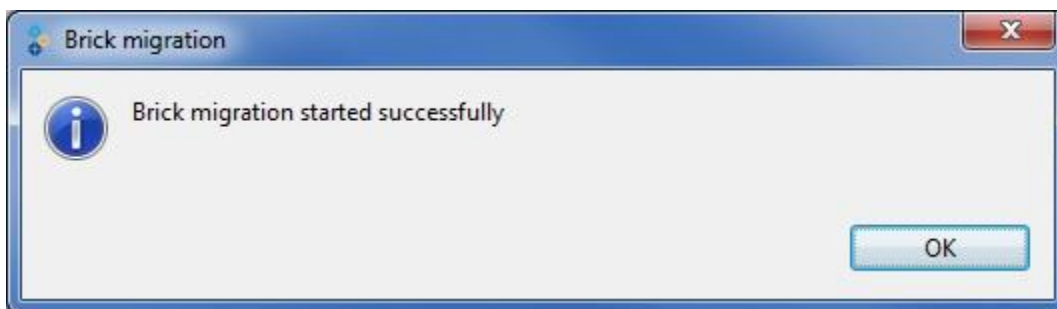
To migrate volume data from one brick to another, perform the following steps.

1. Choose from the menu, toolbar, or right click and select the Volume and select **Migrate Bricks**.

The **Migrate Brick** window is displayed.



2. Select the source brick from the **From Brick** table.
3. Select the destination brick from the **To Brick** table.
4. Click **Finish**. The **Brick Migration** window appears indicating that the brick was successfully migrated.



5. The **Tasks** tab displays the migration status (updated periodically).

10. Managing Logs

You can download volume logs to your local machine to analyze and archive or you can retrieve volume logs based on a search criteria. This section describes how you can manage logs on your storage cluster.

Typically, the log file includes:

- Installation messages
- Removal messages
- Configuration messages
- Informational messages
- Error messages

The Gluster Management Gateway log files are stored in the following locations on your server:

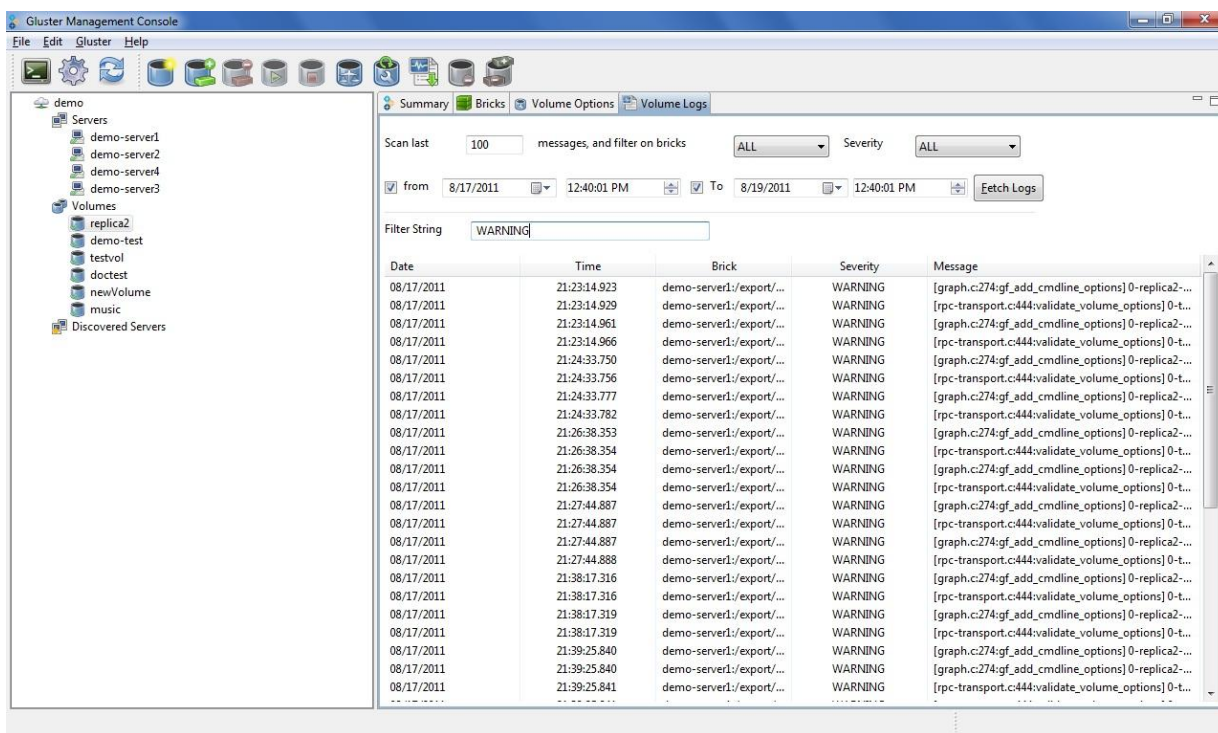
- Tomcat servlet container logs - `/var/log/tomcat5/catalina.out`
- Gluster Management Gateway logs - `/var/log/glustermg/glustermg.log`

10.1. Displaying Log Based on a Search Criteria

To display log based on a search criteria, perform the following steps.

1. Select the volume for which you want to display the log and click the **Volume Logs** tab.
2. Type the number of messages to scan on each brick in the **Scan Last** field.
3. Select **All** or specify bricks from the drop down list.
4. Select the **From date**, **Time**, and **To date** and click **Fetch Logs**.

The volume log messages for the selected criteria is displayed.



The logs fetched from the servers can be filtered further by typing a value in the **Filter String** field.

10.2. Downloading the Logs

You can download the server logs to your local machine to analyze and archive, as required.

To download logs, perform the following steps.

1. Select the volume for which you want to download the logs.
2. Select **Download Logs** from the **Gluster** menu or toolbar.
3. A window is displayed allowing you to choose the download location on your machine.
4. Choose the location, and click **Save**. The logs for the selected volume are downloaded to the location you specified in a `.tar.gz` file format.