Ahsay Offsite Backup Server

v6

High Availability Option Setup Guide

Ahsay Systems Corporation Limited

6 November 2012

Ahsay Offsite Backup Server

High Availability Option Setup Guide

Copyright Notice

© 2012 Ahsay Systems Corporation Limited All rights reserved.
The use and copying of this product is subject to a license agreement. Any other use is prohibited. No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system or translated into any language in any form by any means without priori written consent of Ahsay Systems Corporation Limited Information in this manual is subject to change without notice and does not represent a commitment on the part of the vendor, Ahsay Systems Corporation Limited does not warrant that this document is error free. If you find any errors in this document, please report to Ahsay Systems Corporation Limited in writing.

This product includes software developed by the Apache Software Foundation (http://www.apache.org/).

Trademarks

Ahsay, Ahsay Online Backup Suite, Ahsay Offsite Backup Server, Ahsay Online Backup Manager, Ahsay A-Click Backup, Ahsay Replication Server, Ahsay BackupBox Firmware, Ahsay Universal Backup System, Ahsay NAS Client Utility are trademarks of Ahsay Systems Corporation Limited.

Microsoft, Windows, Microsoft Exchange Server and Microsoft SQL Server are registered trademarks of Microsoft Corporation.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners..

Oracle, Oracle 8i, Oracle 9i, Oracle 10g, Oracle 11g are registered trademarks of Oracle Corporation.

Lotus, Domino, Notes are registered trademark of IBM Corporation.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo and JBoss are registered trademarks of Red Hat, Inc. www.redhat.com in the U.S. and other countries. Linux is a registered trademark of Linus Torvalds.

Apple and Mac OS X are registered trademarks of Apple Computer, Inc.

All other product names are registered trademarks of their respective owners.

Disclaimer

Ahsay Systems Corporation Limited will not have or accept any liability, obligation or responsibility whatsoever for any loss, destruction or damage (including without limitation consequential loss, destruction or damage) however arising from or in respect of any use or misuse of reliance on this document. By reading and following the instructions in this document, you agree to accept unconditionally the terms of this Disclaimer and as they may be revised and/or amended from time to time by Ahsay Systems Corporation Limited without prior notice to you.

Revision History

Date	Descriptions
13 October 2010	First Version of High Availability Guide
26 October 2010	Revise Section 1.4.4
30 June 2011	Updated: 1.1 Conventions 1.2 Definitions, Acronyms and Abbreviations 1.3 System Requirements
6 November 2012	Content Update

Date	Descriptions	Type of modification
10 May 2012	3.3 Setup first AhsayOBS instance in Windows Platform	Modified
11 May 2012	8 Upgrade AhsayOBS	New
25 May 2012	3.4 Configure AhsayOBS as a generic service	Modified
04 Oct 2012	Copyright Notice	Modified

Table of Contents

1 Overview 1
1.1 Conventions 1
1.2 Definitions, Acronyms and Abbreviations 2
1.3 System Requirements 3
1.4 High Availability 3
1.4.1 Failover Cluster Concepts 3
1.4.2 Failover Clustering in Windows Server 2008 6
1.4.3 Failover Clustering in Red Hat Enterprise Linux 5 AP 8
1.5 High Availability AhsayOBS 10
1.5.1 Important Notes 10
1.5.2 HA Scenario on Windows Platform 11
1.5.3 HA Scenario on Linux Platform 14
2 Hardware Setup and Configuration 16
3 Setup High Availability AhsayOBS on Windows 18
3.1 Operating System and Storage Configuration 18
3.2 Failover Cluster Configuration 21
3.3 Setup first AhsayOBS instance in Windows Platform 32
3.4 Configure AhsayOBS as a generic service 34
3.5 Setup additional AhsayOBS in the cluster 44
4 Verify AhsayOBS HA Setup on Windows 45
4.1 Cluster Setup 45
4.2 Cluster Failover 45
4.3 Other Verification Tests 46
5 Setup High Availability AhsayOBS on Linux 47
5.1 Operating system and storage configuration 47
5.2 Failover cluster configuration 52
5.3 Setup first AhsayOBS instance in Linux Platform 66
5.4 Configure AhsayOBS as a high availability service 69
5.5 Setup additional AhsayOBS in the cluster 73
6 Verify AhsayOBS HA Setup on Linux 75
6.1 Cluster Failover 75
6.2 Other Verification Tests 76
6.3 Troubleshooting 76
6.4 Additional Information 77
7 Server Tuning Recommendations 78
8 Upgrade AhsayOBS 80
8.1 Windows environment 80
9 References 83
Appendix 84
Appendix A Product Documentations 84
Appendix B Sample Script for AhsayOBS HA on Linux 85

1 Overview

Ahsay™ Offsite Backup Server (AhsayOBS) is the central core server that accepts and handles backup data from multiple backup clients simultaneously. Backup clients, Ahsay™ Online Backup Manager (AhsayOBM) and Ahsay™ A-Click Backup (AhsayACB) cannot backup data if the core server is being interrupted by service restart or temporarily network disconnection.

The High Availability Setup Option offers the solution to eliminate the single point of failure. AhsayOBS High Availability Option can be achieved by using the Failover Clustering feature on Microsoft Windows Server 2008 / Red Hat Enterprise Linux 5 Advance Platform and selective hardware.

This setup guide will focus on setting up AhsayOBS on a Two-Node Failover Cluster that comprises of 2 Dell PowerEdge R710 Servers with 1 Dell PowerVault MD3000 disk array system. The cluster can be expanded according to the system scale.

1.1 Conventions

Convention	Descriptions	Example
Bold	Important Information	**Important: Please do not activate the license before AhsayOBS is configured as a generic service in windows 2008.**
Italic	Folder Path or File Path	*C:\Program Files\AhsayOBS*
[]	Graphical Interface Elements	[Manage Log]
%%	Filepath in Windows format	%OBSR_HOME%
$	Filepath in Linux format	$OBSR_HOME
Italic	Command	*mkdir /usr/local/obs*

Notation that are used specifically for API section:

Convention	Descriptions
Angle brackets <Text>	Required parameter
Square brackets [Text]	Optional parameter
Braces {Text}	Set of required value; select one

Vertical bar (|)	Separator for mutually exclusive items; select one
Ellipsis (…)	Repeat-able values

1.2 Definitions, Acronyms and Abbreviations

Term / Abbreviation	Definition
AhsayOBS	Ahsay™ Offsite Backup Server
AhsayRDR	Ahsay™ Redirector Server
AhsayOBM	Ahsay™ Online Backup Manager
AhsayACB	Ahsay™ A-Click Backup
OBSR_HOME	The install location of AhsayOBS & AhsayRPS Windows *C:\Program Files\AhsayOBS and AhsayRPS* Linux /usr/local/obs/
OBSR_HOME_BAK	The backup location of AhsayOBS & AhsayRPS Windows *C:\Program Files\AhsayOBS and AhsayRPS.bak* Linux */usr/local/obs.bak/*
USER_HOMES	Directory where backup accounts Are stored
SYSTEM_HOME	The directory where AhsayOBS and AhsayRPS store its system information. Windows %OBSR_HOME%\system\ Linux $OBSR_HOME\system\
TCP/IP	Transmission Control Protocol / Internet Protocol.
Client Access Point	The IP assign for a generic service in a cluster setup.
AP	This is the abbreviation of Advance Platform

1.3 System Requirements

To setup the AhsayOBS with High Availability option, please ensure the following requirements are met:

Hardware

1. At least 2 servers that meet our recommended hardware configuration. Please refer to Ahsay Help Centre Article #2323 for further details.

2. At least one Disk Array System designed for cluster operation.

Software

1. The same version of operating system is installed on all servers involved in the HA setup. Please use one of the following OS for the HA Setup:

    - Windows Server 2008 Enterprise / Datacenter Edition
    - Windows Server 2008 R2 Enterprise / Datacenter Edition
    - Red Hat Enterprise Linux 5.3 Advance Platform

2. The AhsayOBS License key used in the HA setup must have the HA server add-on module enabled.

1.4 High Availability

High Availability is a system design approach to minimize the downtime during a system breakdown. To provide high availability web services, failover cluster is the most common setup.

1.4.1 Failover Cluster Concepts

Failover cluster is also known as high availability cluster. It is an implementation of computer cluster that is designed primarily for providing high availability service. This involves redundant hardware and software, e.g. computers and storage system. With a failover cluster, the service downtime will be minimized when a breakdown occurs. This is because the cluster is able to detect the breakdown automatically and restart the service on another redundant computer.

In a failover cluster setup, each computer involved is called a node and usually they are all connected with a private network connection, known as heartbeat. This private network connection is used to monitor the health and status of each node.

There are different failover cluster sizes and the most common one is a Two-Node Failover Cluster. This is because it requires the minimum resource to build.

The Two-Node Failover Cluster consists of two servers and usually it is connected to a shared storage. There are two main type of model for Two-Node Failover Cluster configuration:

Active/Active Clustering Configuration

Each node is running its own service while providing failover capabilities for resource of the other node.

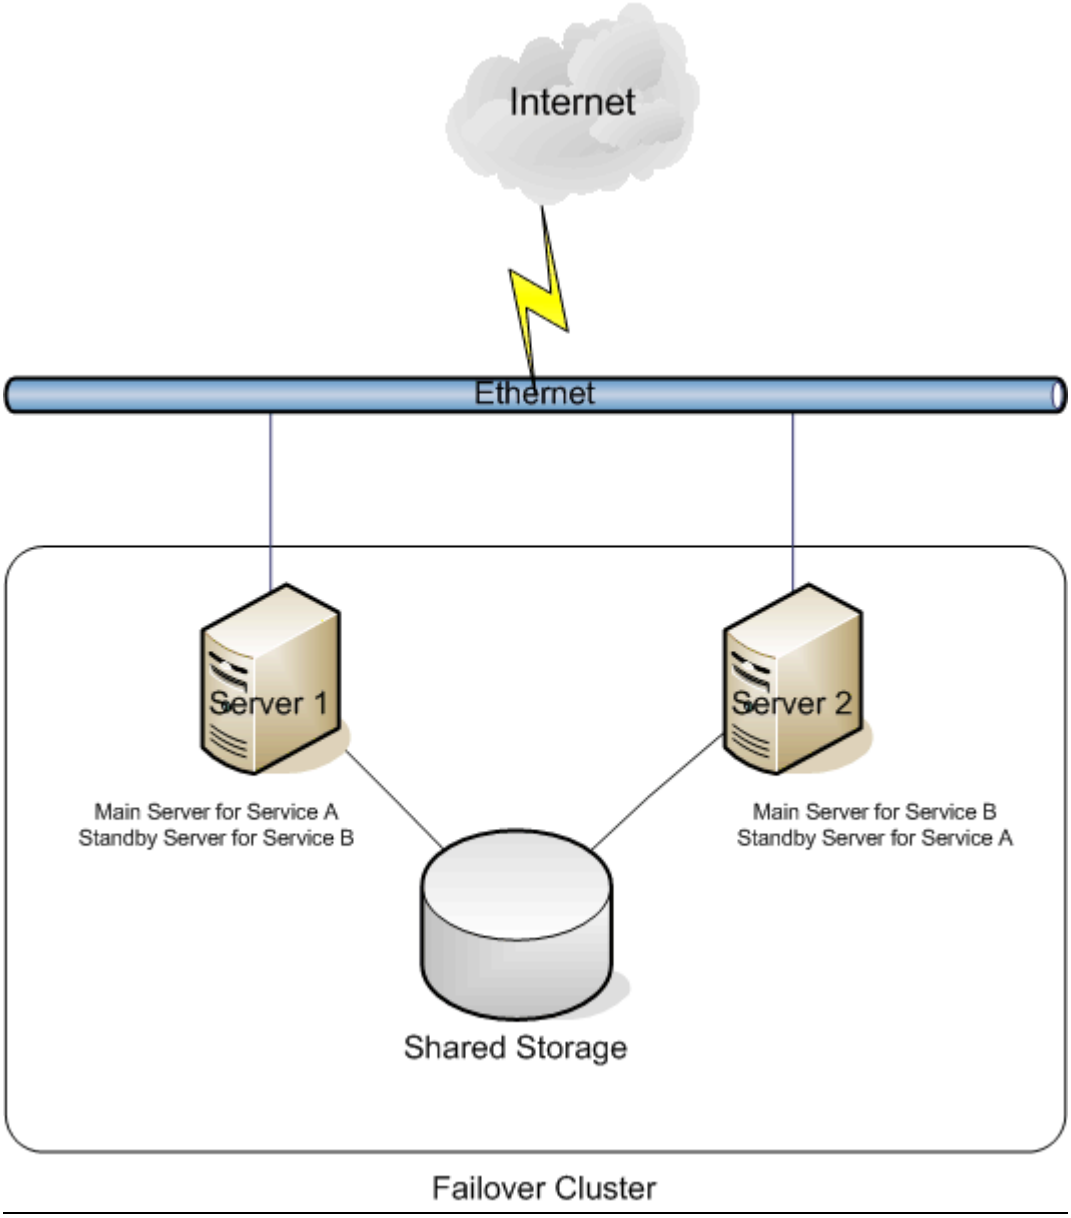

Active/Passive Clustering Configuration

Only one node provides the service while the other is only brought online when the primary node fails.

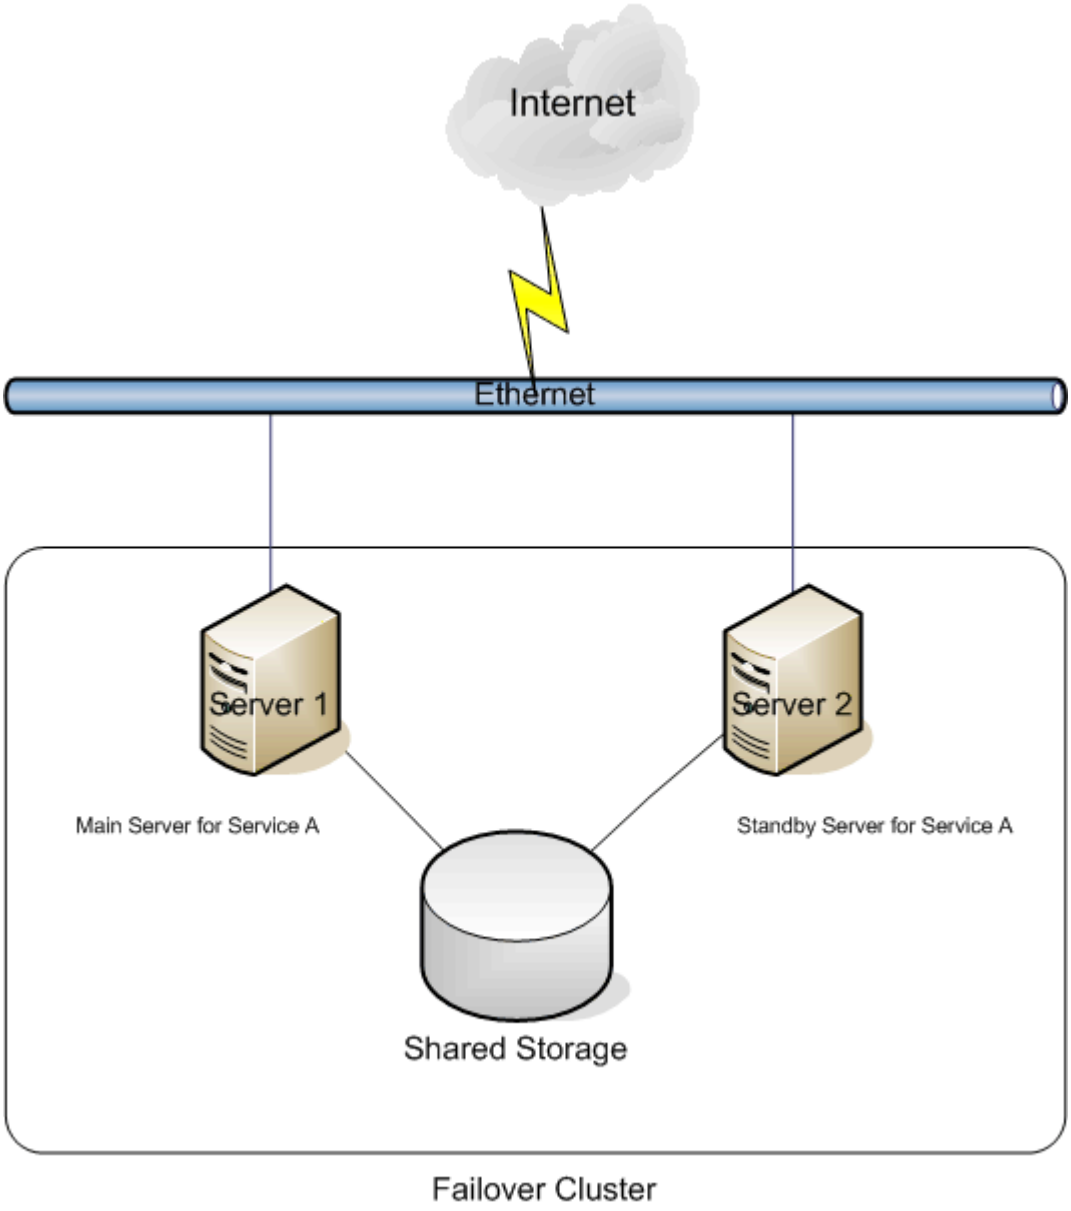

1.4.2 Failover Clustering in Windows Server 2008

The Failover Clustering feature of Windows Server 2008 is used to setup and manage the behavior of the failover cluster. It is an add-on feature available in Windows Server 2008 Enterprise Edition or Datacenter Edition.

The following are some technical terms related to Failover Clustering:

Quorum

In Windows 2008, the quorum of the cluster is the number of elements that must be online for the cluster to work properly. Under normal circumstances, this information should be accessible to all nodes in the cluster.

Generic Service

This is a type of service that is run on the cluster node and it is managed by the cluster software. The cluster software will start the service on one of the nodes in the cluster, and check the status of the service periodically to determine whether the service appears to be running.

The generic service is comprised three major components, e.g. IP address, storage and the service itself.

During a failover, those resources will be detached from the original node and reattach to the destination node, which will then bring the service online with exactly the same set of resources.

Failover

Failover occurs when the current service node is down and the cluster software will choose another available node to continue providing the generic service.

Service Node

Service Node is the node in the cluster which is providing the generic service.

A node can become the service node (e.g. run a generic service) if it possesses all its associate resources, therefore the possible owner of a generic service is one that can satisfy this condition. For example, consider the below dependency diagram:

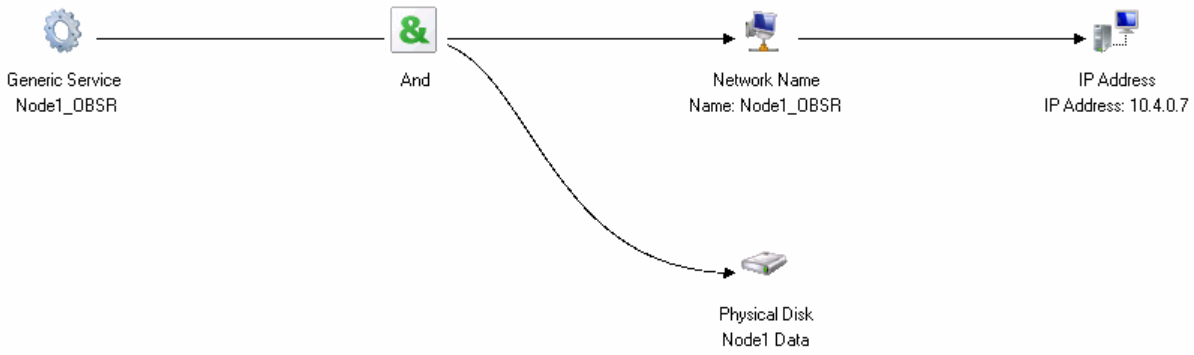

Node 1 is a possible owner of Generic Service Node1_OBSR if it is a possible owner of all elements in the diagram.

Before a generic service is able to failover to a node, it must first satisfy two conditions:

1. The Node must be a possible owner of the generic service.

2. The service must be registered at the node.

Two-Node Clustering

Two-Node Clustering Model (also known as Node and Disk Majority Model) is one of the available cluster models in Windows Server 2008. This model requires at least 2 nodes and 1 shared quorum. It can run properly under one of the following conditions:

- The shared quorum is available but half of the nodes are down.
- The shared quorum is not available with over half of the nodes are available.

1.4.3 Failover Clustering in Red Hat Enterprise Linux 5 AP

The Red Hat Cluster Suite of Red Hat Enterprise Linux 5 AP is used to setup and manage the behavior of the failover cluster. It is an advance feature available in Red Hat Enterprise Linux 5 Advance Platform Edition. This setup is only recommended for experienced users in Linux platform.

The following are some technical terms related to the Red Hat Cluster Suite:

cluster.conf

The *cluster.conf* contains the cluster configuration per-clustered Node, this configuration is placed under /etc/cluster folder of every Node. User should NOT configure this file manually unless instructed. Red Hat has provided a graphical interface in modifying the cluster configuration, such that any changes will be propagated to all cluster nodes automatically.

Conga

Conga is a web base cluster administrative tools provided by Red Hat, it is an integrated set of software components that provides centralized configuration and management of clusters and storage. It operates on a server/agent model, where the server is known as luci and agent known as ricci. Luci communicates to the client agent on nodes to obtain information and manages its cluster configurations.

CMAN

CMAN is the abbreviation for Cluster Manager, it is distributed across all nodes in the cluster. CMAN manages the cluster quorum and cluster membership, in another words it determines whether the node is “alive” and whether the node is a member of the cluster.

Quorum

In Red Hat Cluster Suite, the quorum of the cluster is the number of votes CMAN must have for the node to be granted membership, it is usually (n/2 + 1).

High Availability Service

A high availability service is a service that runs on the cluster node and is managed by the high availability service manager (rgmanager). rgmanager will start the service on one of the nodes in the cluster, and check the status of the service periodically to determine whether the service appears to be running.

The high availability service is comprised numbers cluster resources, e.g. IP address, storage or an application initialization script.

During a failover, those resources will be “detached” from the original node and “reattach” to the destination node, which will then bring the service online with exactly the same set of resources.

Failover

Failover occurs when the current service node is down and the cluster software will choose another available node to continue providing the High availability service.

Two-Node Clustering

Two-Node Clustering is a special occasion, since to become quorum at least 1 vote is required, which is not practical. In order for a two-node cluster to operate, it requires a shared quorum. It can run properly under one of the following conditions:

- The shared quorum is available but half of the nodes are down.
- The shared quorum is not available with over half of the nodes are available.

1.5 High Availability AhsayOBS

To provide a high availability AhsayOBS, the AhsayOBS should be configured as a generic service in a Windows 2008 Failover Cluster or a high availability service in Red Hat Cluster Suite.

The backup data must be stored in a disk array system with data protection feature.

With this configuration, the AhsayOBS service and all its required resources will be moved to another available node when the service node is down, and hence, the disruption to the backup service is minimized.

1.5.1 Important Notes

Please ensure that the following requirements are met before setting up AhsayOBS High Availability:

Windows Server 2008

1. All cluster nodes must belong to a common domain. For best practice, the domain controllers should not be used as clusters nodes.

2. All cluster nodes are on the same edition of Windows 2008 / 2008 R2.

3. A minimum of 2 nodes are required.

4. For Active / Active configuration, please ensure that each node can sustain the load of running two instances of AhsayOBS.

5. Each AhsayOBS instance must hold a valid license with High Availability server add-on module enabled.

Important:
Please do not activate the license before AhsayOBS is configured as a generic service in windows 2008.

Red Hat Enterprise Linux 5 AP

1. All cluster nodes are on the same edition of Red Hat Enterprise Linux 5 AP.

2. All cluster nodes are using the same version of Red Hat Cluster Suite components.

3. A minimum of 2 nodes are required.

4. For Active / Active configuration, please ensure that each node can sustain the load of running two instances of AhsayOBS.

5. Each AhsayOBS instance must hold a valid license with High Availability server add-on module enabled.

Important:
Please do not activate the license before AhsayOBS is configured as a high-availability service in Red Hat Enterprise Linux.

1.5.2 HA Scenario on Windows Platform

The following scenario is used as an example to explain the steps required to configure High Availability AhsayOBS on the Windows platform.

In this example, we are going to setup an Active / Active Two-Node Failover Cluster with two Dell PowerEdge R710 Servers with a shared Dell PowerVault MD3000 Disk Array System.

A diagram about the planned setup of this scenario is shown below:

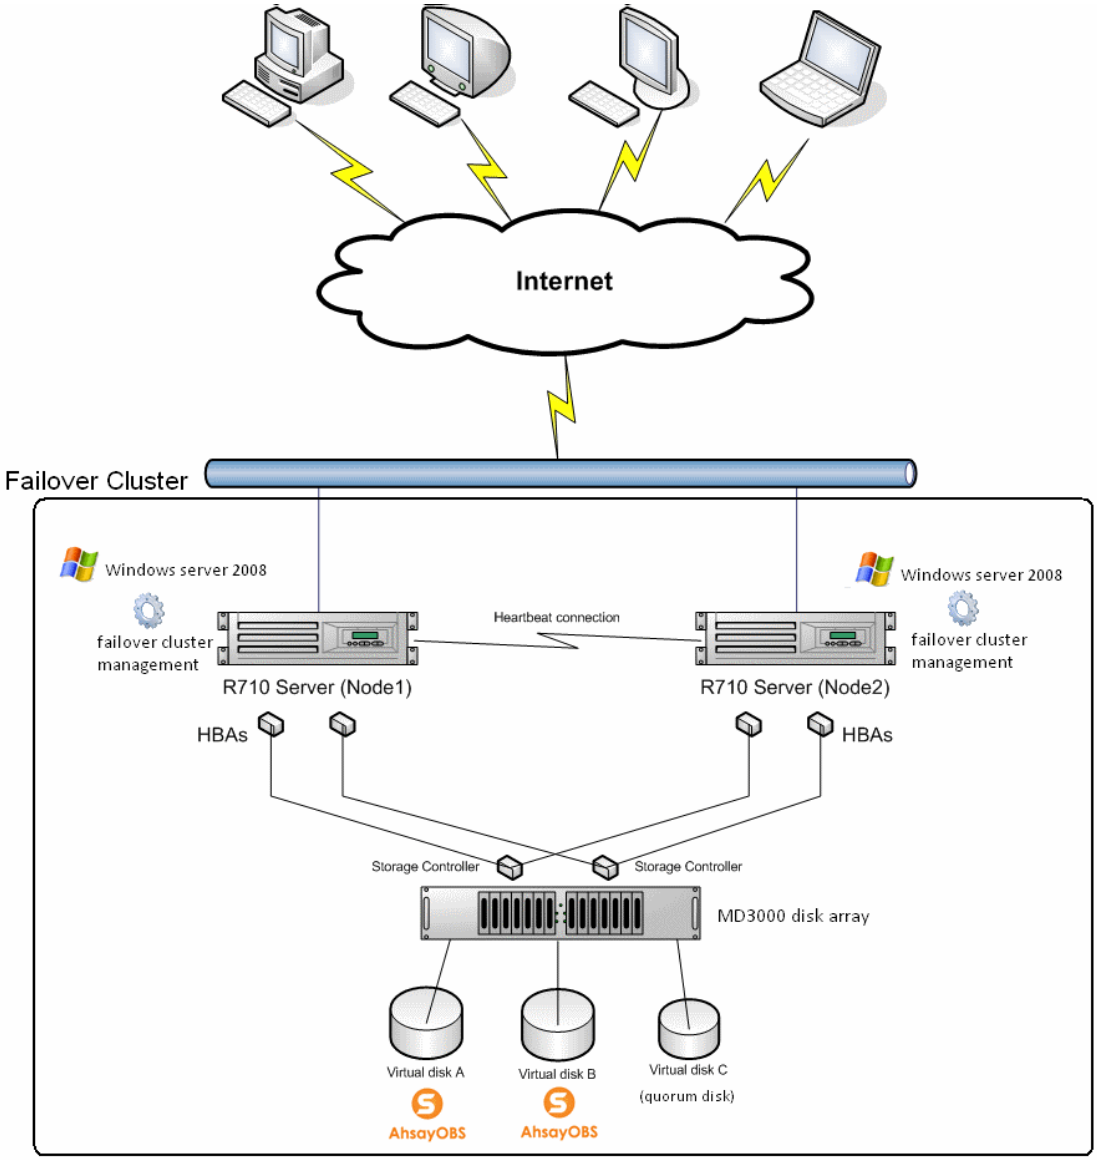

The information of the 2 servers is listed below:

	Server 1	Server 2
Computer Name	Node1	Node2
IP Address	10.4.0.4	10.4.0.5
Cluster Name	HA_Clustering	
Cluster IP Address	10.4.0.6	

The information of AhsayOBS installed on both nodes:

Server	Service IP	Installation Location	Ports Used
Node1	0.0.0.0	G:\AhsayOBS and AhsayRPS	80, 443 & 8014
Node2	0.0.0.0	H:\AhsayOBS and AhsayRPS	70, 8443 & 8015

The information of generic services of the cluster is listed below:

	Generic Service 1	Generic Service 2
Generic Service Name	Node1OBSR	Node2OBSR
Service IP Address	10.4.0.7	10.4.0.8

Note:
In this documentation, the cluster nodes are installed on domain controllers for testing purposes. However, in product environment the domain controllers should never be used as cluster nodes.

1.5.3 HA Scenario on Linux Platform

The following scenario is used as an example to explain the steps required to configure High Availability AhsayOBS on the Linux platform.

In this example, we are going to setup an Active/Active Two-Node Failover Cluster with two Dell PowerEdge R710 Servers with a shared Dell PowerVault MD3000 Disk Array System.

A diagram about the planned setup of this scenario is shown below:

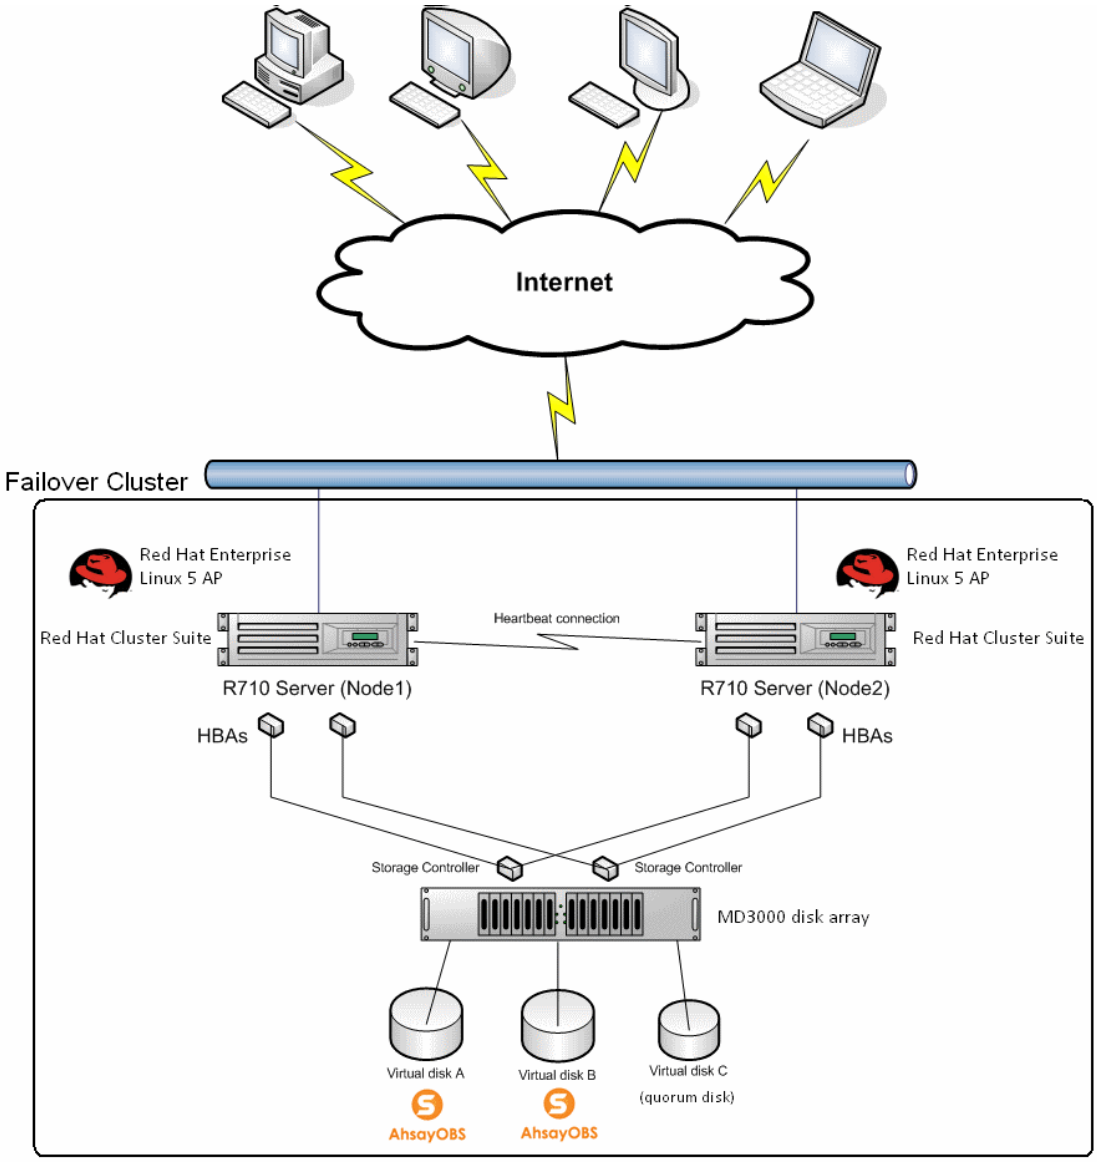

The information of the 2 servers is listed below:

	Server 1	Server 2
Computer Name	Node1	Node2
IP Address	10.4.0.4	10.4.0.5
Cluster Name	HACluster	

The information of AhsayOBS installed on both nodes:

Server	Service IP	Installation Location	mount device	Ports Used
Node1	0.0.0.0	/mnt/mount1/OBSR1	/dev/VG1/LV1	80, 443 & 8014
Node2	0.0.0.0	/mnt/mount2/OBSR2	/dev/VG2/LV2	70, 8443 & 8015

The information of High Availability Service of the cluster is listed below:

	HA Service 1	HA Service 2
HA Service Name	Node1OBSR	Node2OBSR
Service IP Address	10.4.0.7	10.4.0.8

2 Hardware Setup and Configuration

This chapter includes steps of basic hardware configuration required in setting up a Two-Node Failover Cluster with two Dell PowerEdge R710 and a single Dell PowerVault MD3000 Storage Array, as shown in the first diagram of this document. For more information on the hardware configuration, please refer to the documentations listed in the References section.

Please follow the steps mentioned below to configure your hardware:

Step 1 - Cabling the Power supplies

- For nodes with multiple power supplies, plug each power supply into a separate AC circuit.
- Use uninterruptible power supplies.

Step 2 - Cabling your public and private networks

The network adapters in R710 Server provided at least two network connection, one for the public and one for the private network.

1. Connect the public network to a network adapter that supports TCP/IP.
2. Connect the two R710 servers through the private network adapter, this network is used for intra-cluster communications.

Below diagram shows such network cabling configuration:

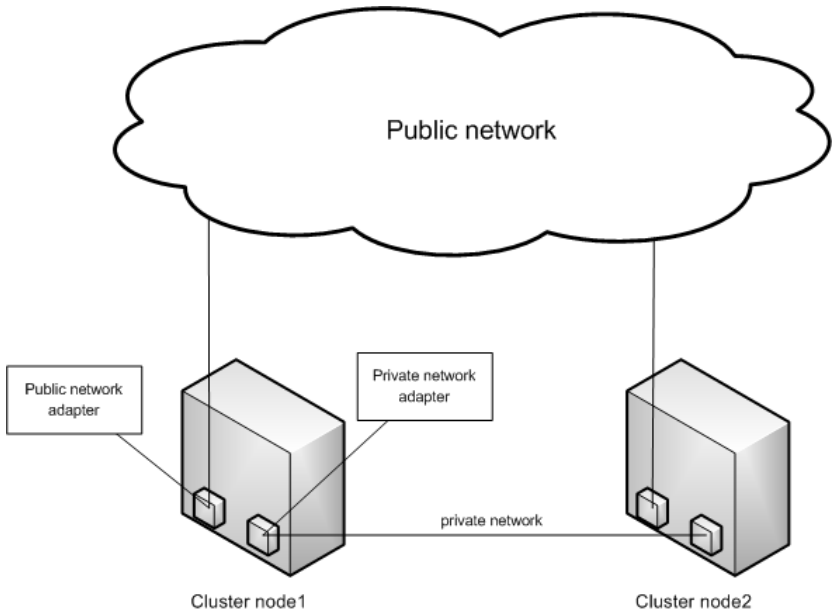

Step 3 - Cabling your storage system in redundant configuration with Dual SAS 5/E HBAs

1. Connect cluster Node1 to MD3000,

   i. Install a SAS cable from cluster node 1 HBA 1 port 0 to the RAID controller MD3000 module 0 port In-0.

   ii. Install a SAS cable from cluster node 1 HBA 2 port 0 to the RAID controller MD3000 module 1 port In-0.

2. Connect cluster Node2 to MD3000,

   iii. Install a SAS cable from cluster node 2 HBA 1 port 0 to the RAID controller MD3000 module 0 port In-1.

   iv. Install a SAS cable from cluster node 2 HBA 2 port 0 to the RAID controller MD3000 module 1 port In-1.

Below shows such network cabling connections:

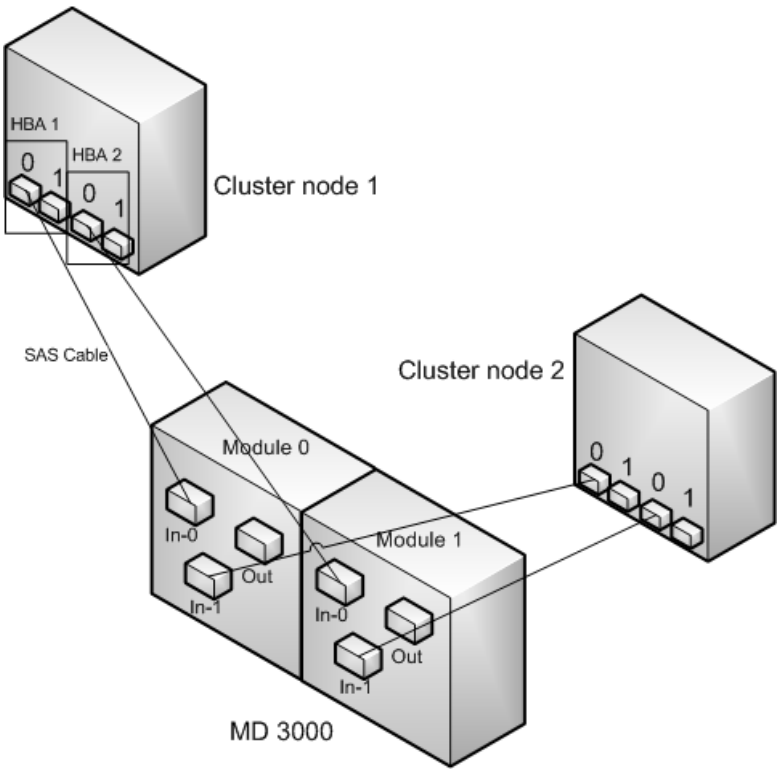

3 Setup High Availability AhsayOBS on Windows

This chapter describes how to setup the AhsayOBS high availability on Windows Platform.

AhsayOBS High Availability setup is summarized as follows:

1. Operating system and storage configuration
2. Failover cluster configuration.
3. Setup the first AhsayOBS instance.
4. Configure AhsayOBS as a generic service.
5. Setup additional AhsayOBS instance(s) to the cluster. (This step can be repeated for as many AhsayOBS Instances for the cluster as required.)

3.1 Operating System and Storage Configuration

When the R710 servers and MD3000 are connected, install Windows Server 2008 onto each of the cluster node. After that, install Dell's Modular Disk Storage Manager onto each of the cluster node, this is used for MD3000 disk array configuration. You will also need to install the SAS 5/E drivers and multi-path driver on each Node, these are available in the MD3000 Resource disk.

Building your virtual disks

The Dell PowerVault MD3000 Disk Array System is comprised 15 individual disks, though independently they can be integrated into a virtual disk and be used as is a single disk. As shown in the diagram in the chapter 1.5, three virtual disks are created, two of which will be used for installing AhsayOBS, and one will be used as a quorum disk, which will be covered later in this document.

Please follow the instructions below to build the virtual disk:

1. Open [Modular Disk Storage Manager].
2. Go to the [Configure] tab and select [Create Disk Groups and Virtual Disks] option.

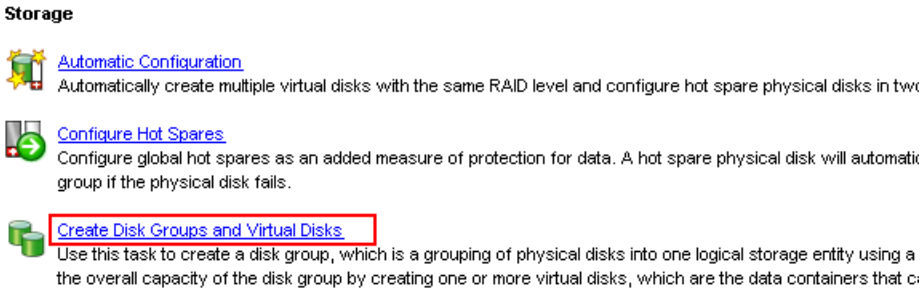

3. Choose [Disk group] and click next.

4. Choose the desire RAID level and number of disk.

5. After that, partition the disk group into virtual disks.

The following settings are used in the Active / Active Two-Node Failover Cluster in our example:

Virtual Disk	Disk Group	RAID Level	Size	Host Group
Virtual Disk A	Disk 0 – 5	RAID5	4.65TB	Host Group 1
Virtual Disk B	Disk 6 – 10	RAID5	3.72TB	Host Group 1
Virtual Disk C	Disk 11 – 12	RAID1	2GB	Host Group 1

Note:
Not all available space has to be used.

Mapping virtual disk to Host Group

Each cluster Node connected to MD3000 is seen as a Host, in order for a Host to access a virtual disk, we need to explicitly map the disk to the Host.

1. Create host group by selecting [Create Host Group] option under [Configure] tab, enter a name (e.g. Host Group 1), add both cluster Node 1 and Node 2 into the Host Group and then click [OK].

2. Select [Create Host-to-Virtual Disk Mappings] option under [Configure] tab, select [Host Group 1] and map Virtual Disk A, B and C to it.

Now all Virtual Disks are mapped to Hosts in Host Group 1, so that both cluster nodes are able to access these virtual disks.

Format the virtual disk

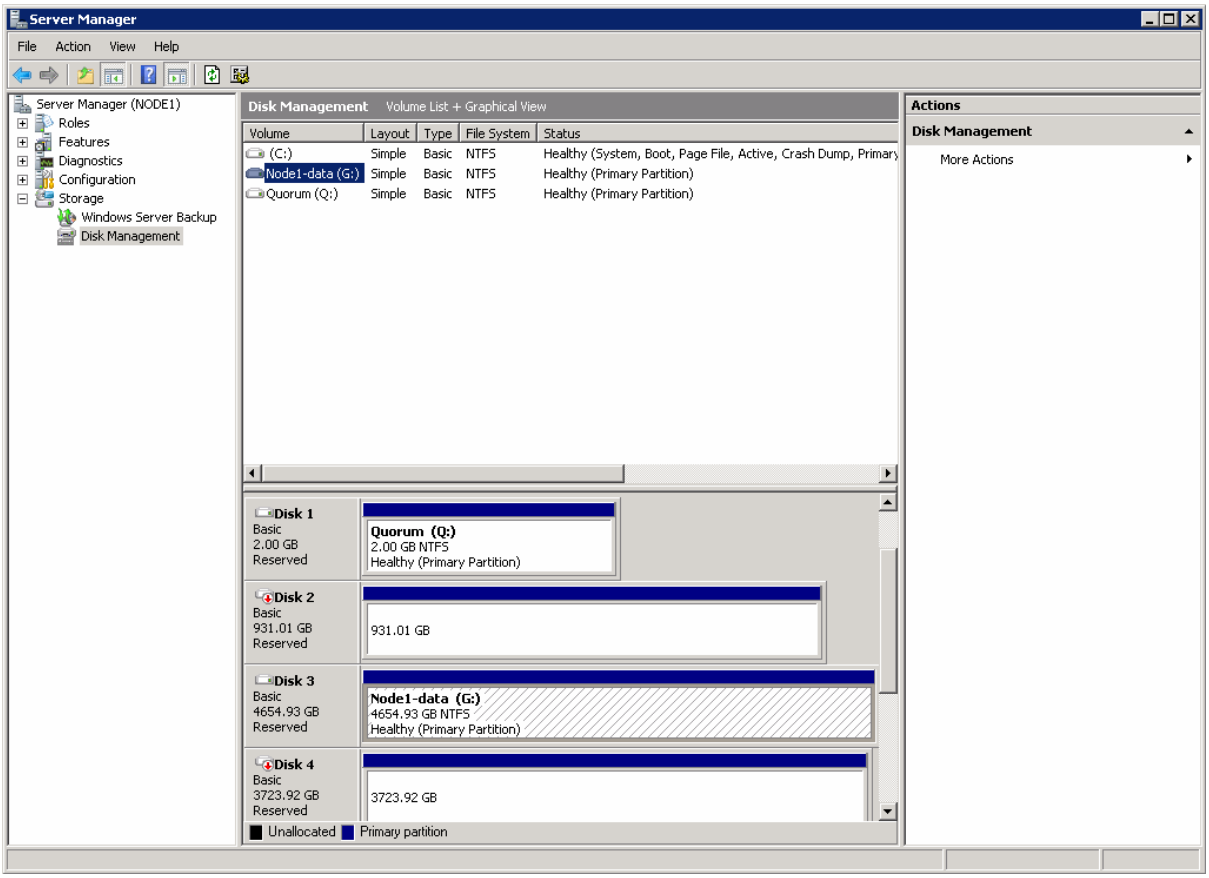

We can now see these virtual disks on Windows installed on both cluster nodes. To format the disks, open [Server Manager] in one of the nodes, go to [Disk Management] under [Storage] section. Apart from the server's local disk, you should use the virtual disks created in step 5. Please format them into NTFS File System.

In our scenario, we are using the following the settings:

Disk Drive	Associate Virtual Disk	File System
Drive G	Virtual A	NTFS
Drive H	Virtual B	NTFS
Drive Q	Virtual C	NTFS

3.2 Failover Cluster Configuration

This chapter focuses on the required setup and configuration on Windows Server 2008 in order to create the Failover Cluster.

Step 1 – Create a Windows Domain

Having completed the hardware configuration, you are now ready to start creating your cluster.

In order to create a cluster, we must first create a windows domain. This is because the [Failover Cluster] feature must be run under a domain environment.

To create a domain, we should first go to Node 1 and perform the following steps:

1. Click [Start] > [Run] then type *dcpromo* and press [Enter]. The active Directory Domain Service Installation Wizard will prompt.

2. Follow the steps shown in the wizard to create a new domain. Remember to choose [Yes] when it asks whether to install DNS Server.

Node 1 becomes the first Active Directory Domain Controller of the new domain. The Domain Controller is responsible to store domain-related data and manage communications between users and servers, e.g. authentication and directory searches.

After we have finished setting up the first Domain Controller, we should promote Node 2 to Domain Controller to ensure that at least one Active Directory Domain Controller is reachable at all times.

Please follow the instructions below to promote Node 2 as a Domain Controller of the domain created in previous steps:

1. Go to [Start Menu] > [Run] and type *dcpromo* and press [Enter].The active Directory Domain Service Installation Wizard will prompt.

2. Follow the steps shown in the wizard to add a new domain controller to an existing domain. Choose [Yes] when it asks whether to install DNS Server.

You can validate the domain configuration as [Administrative tools] > [Active Directory Users and Computers]. You should be able to see both Node 1 and Node 2 inside the Domain Controller folder (Organization Unit).

Step 2 – Creating the Failover cluster

To create a failover cluster, you need to install both [Failover Clustering] and [Multipath I/O] features at all cluster nodes. By default, [Multipath I/O] should be installed with [Active Directory Domain Controller] role.

Perform the following steps on each node to install the [Failover Clustering] feature:

1. Open [Server Manager] > [Feature], click on [Add features]. A wizard will be prompted.

2. Select [Failover Clustering] feature and click [Install] to initiate the installation.

3. Repeat Step 1 and 2, if [Multipath I/O] feature is not installed.

After installing the [Failover Clustering] feature, open the [Failover Cluster Management] interface via [Start Menu] > [Administrative Tools] > [Failover Cluster Management].

The [Failover Cluster Management] interface is a useful tool in managing the failover cluster. To create a new cluster, please follow the steps listed below:

1. Click on the [Create a Cluster] button under the [Actions] section of the right panel, the [Create Cluster Wizard] will appear.

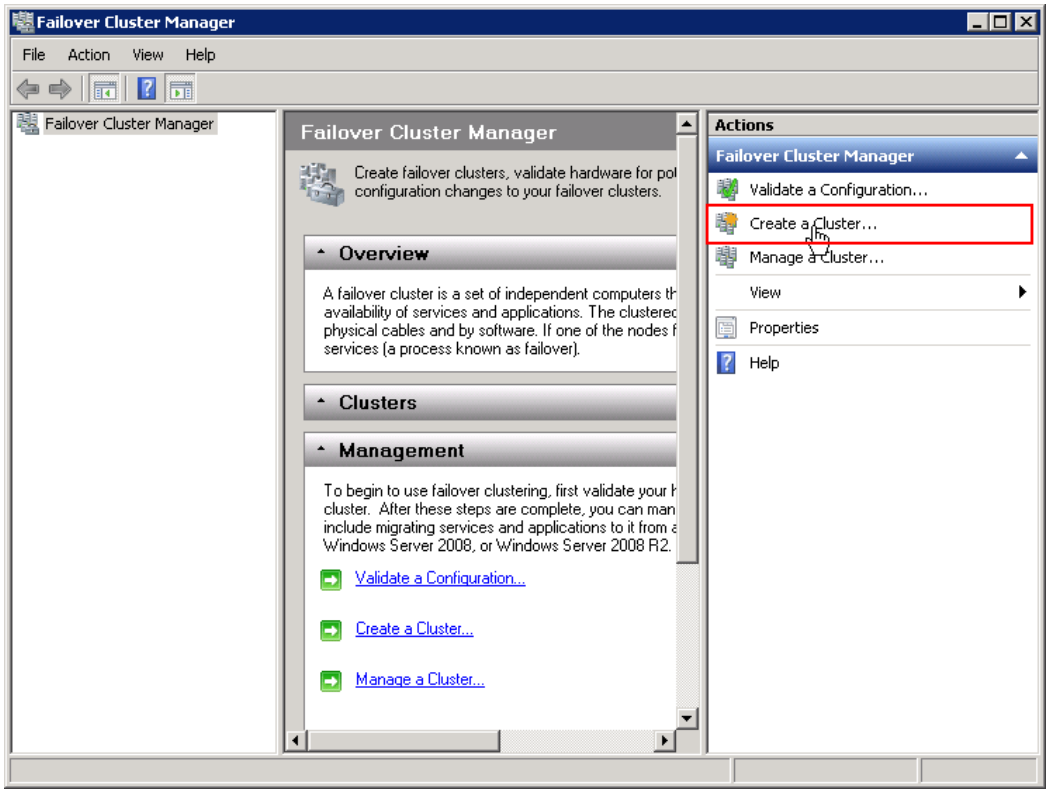
Failover Cluster Manager
File
Action
View
Help
Failover Cluster Manager
Failover Cluster Manager
Create failover clusters, validate hardware for pot
configuration changes to your failover clusters.
Overview
A failover cluster is a set of independent computers th
availability of services and applications. The clustered
physical cables and by software. If one of the nodes f
services (a process known as failover).
Clusters
Management
To begin to use failover clustering, first validate your h
cluster. After these steps are complete, you can man
include migrating services and applications to it from a
Windows Server 2008, or Windows Server 2008 R2.
Validate a Configuration...
Create a Cluster...
Manage a Cluster...
Actions
Failover Cluster Manager
Validate a Configuration...
Create a Cluster...
Manage a Cluster...
View
Properties
Help

2. Click the [Next] button to continue.

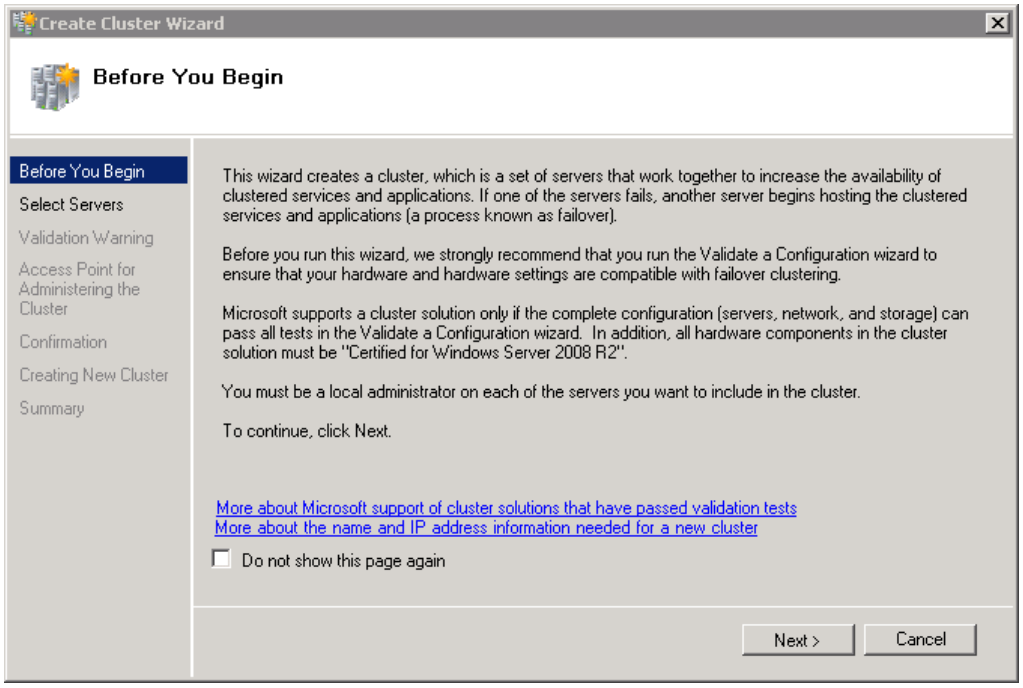

3. Select all servers to be the members of the cluster. In this case, this would be Node 1 and Node 2.

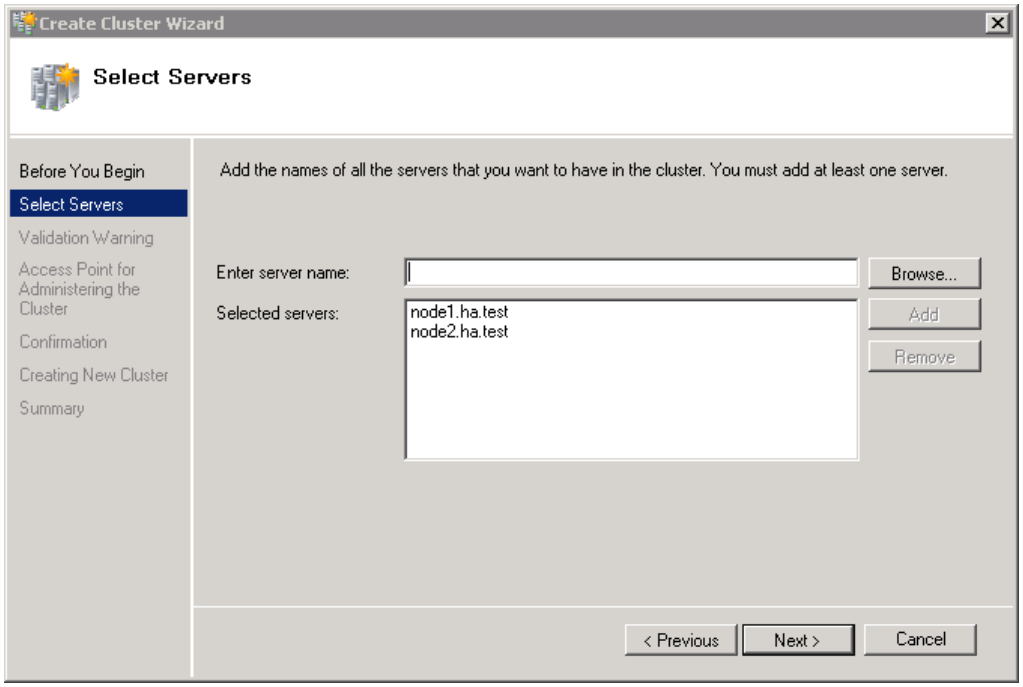

4. If this is the first time for setting up a cluster, it will run the validation tests required.

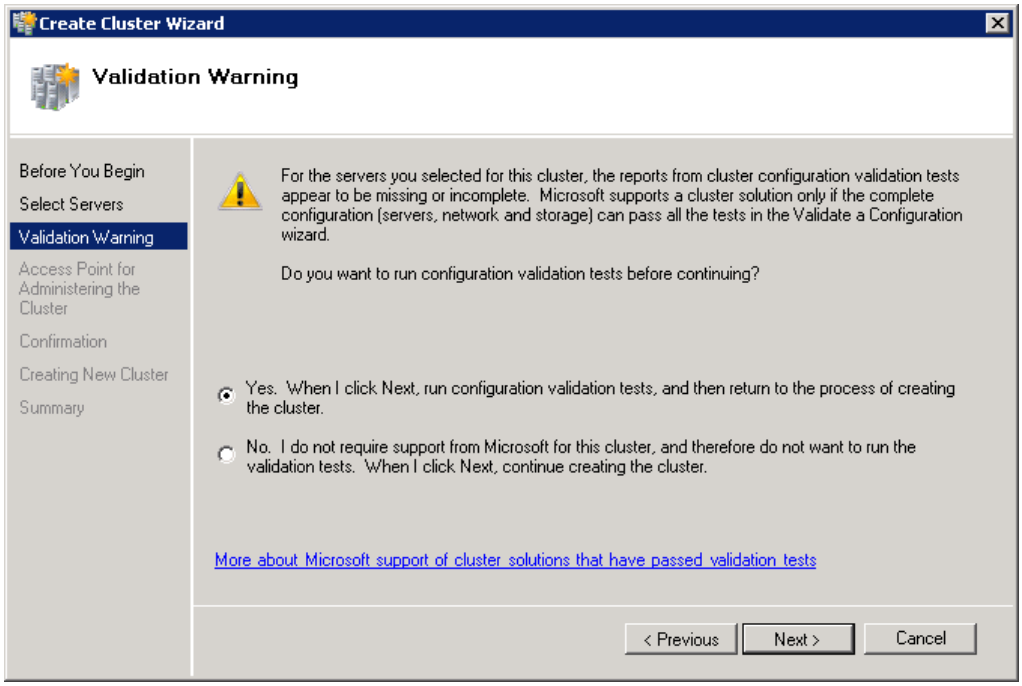

5. On the Validate a Configuration Wizard, click [Next] to go to next step.

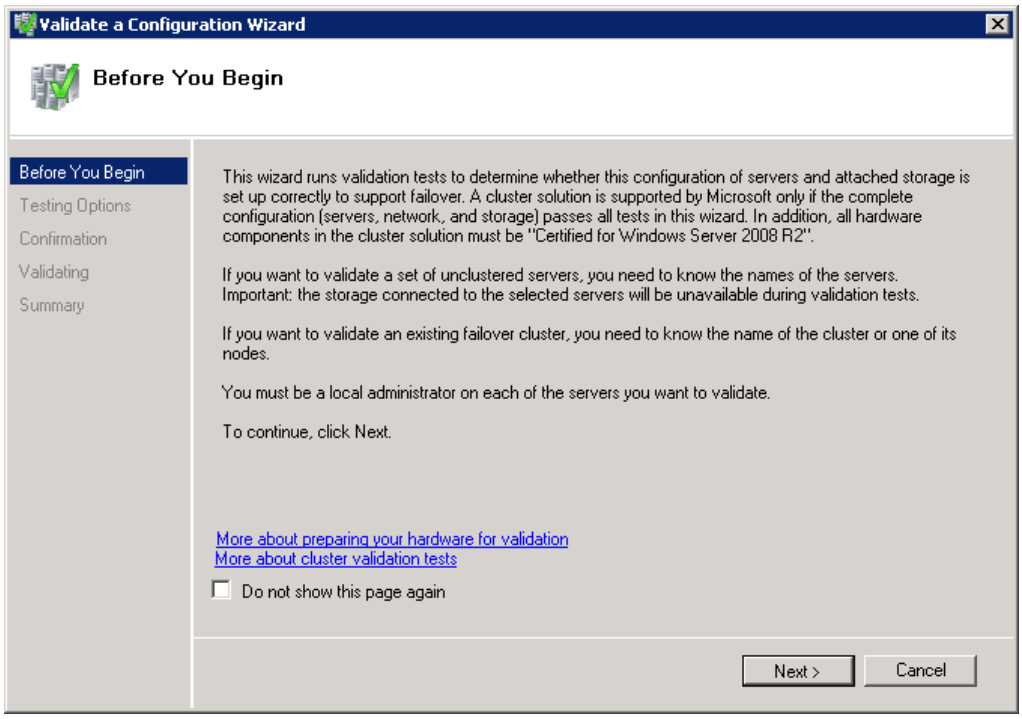

6. Select [Run All Tests].

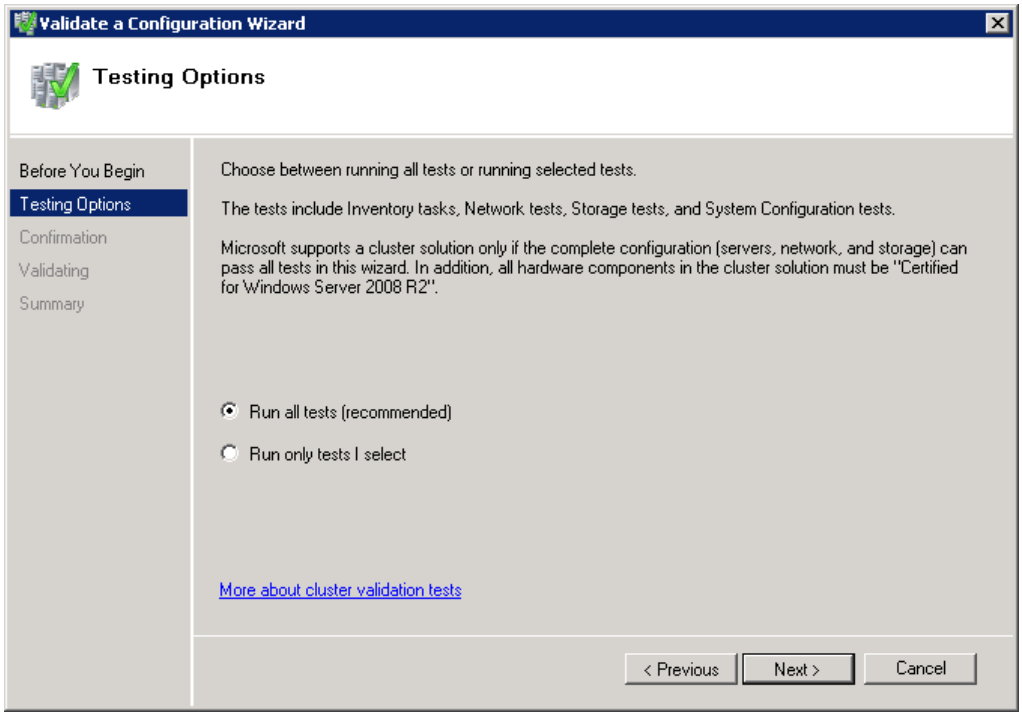

7. The confirmation screen will then be shown on the screen.

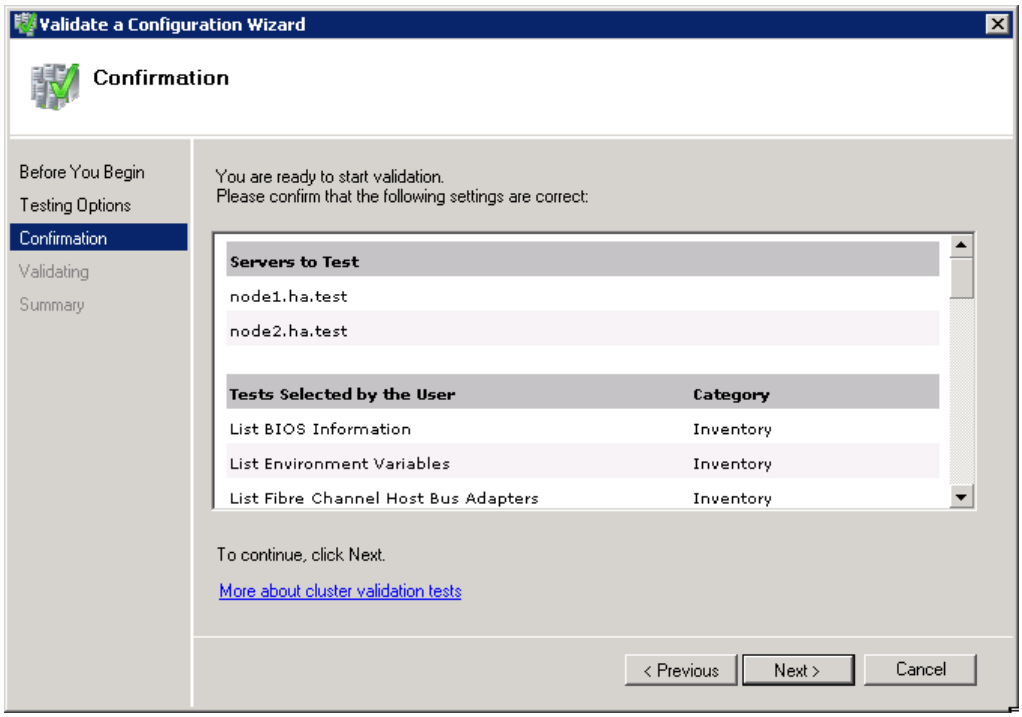

8. Wait until it has finished running all the validation tests.

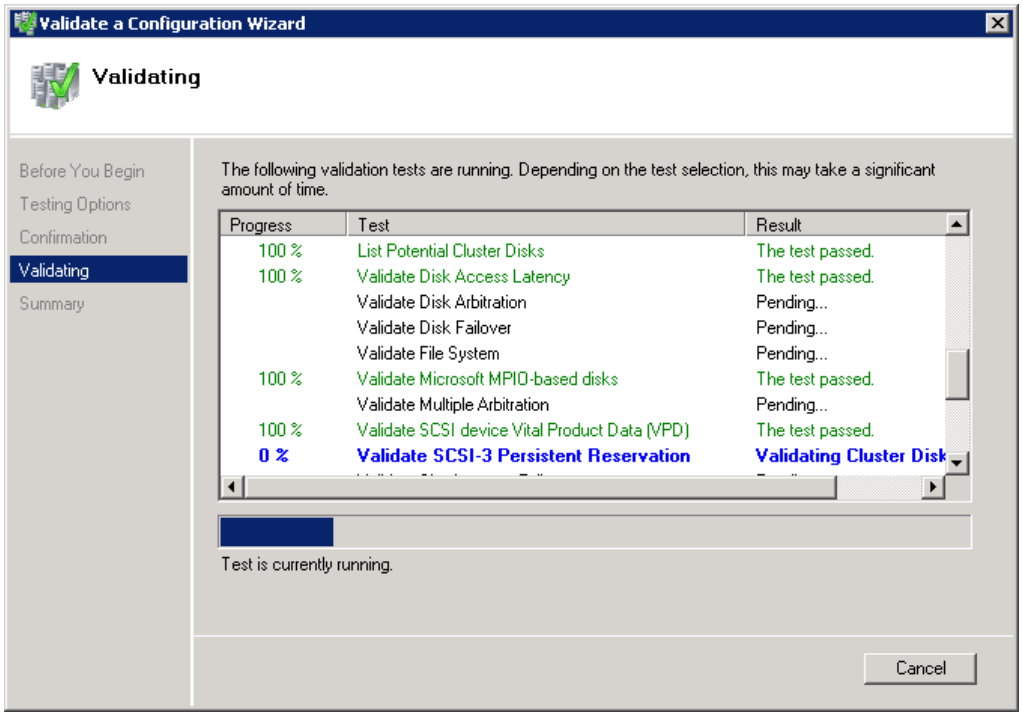

9. A summary report is shown after all the validation tests are completed.

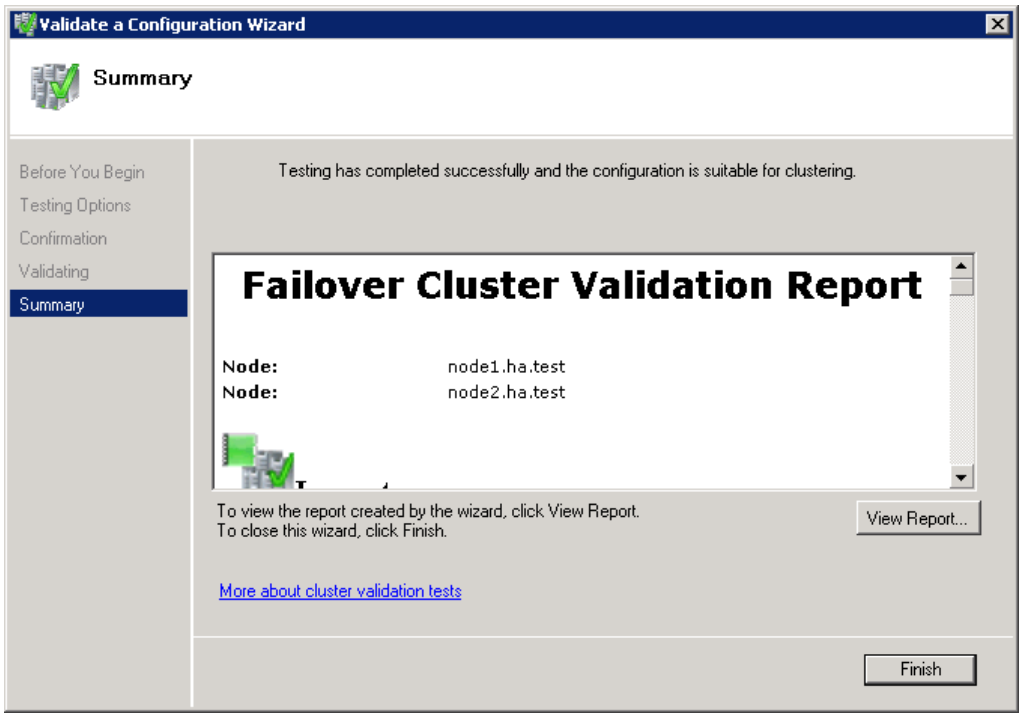

10. Enter a Cluster Name, select the network interface to use and dedicate an IP address for the cluster e.g. 10.4.0.6.

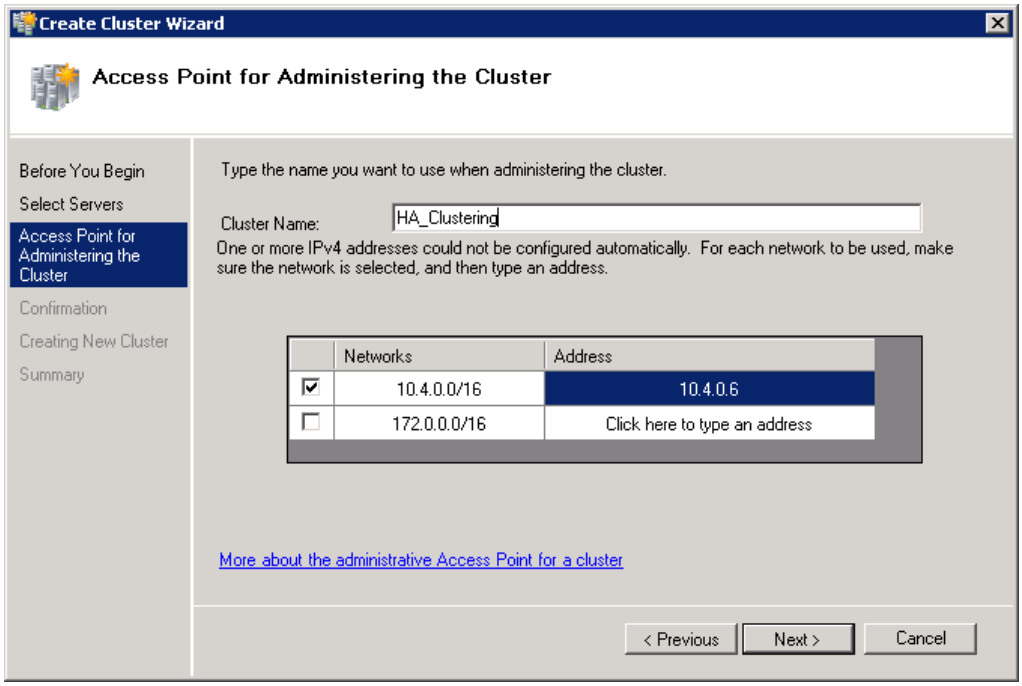

11. If the setup information is correct, click the [Next] button to continue. It may take a few minutes to validate the cluster settings and create the cluster.

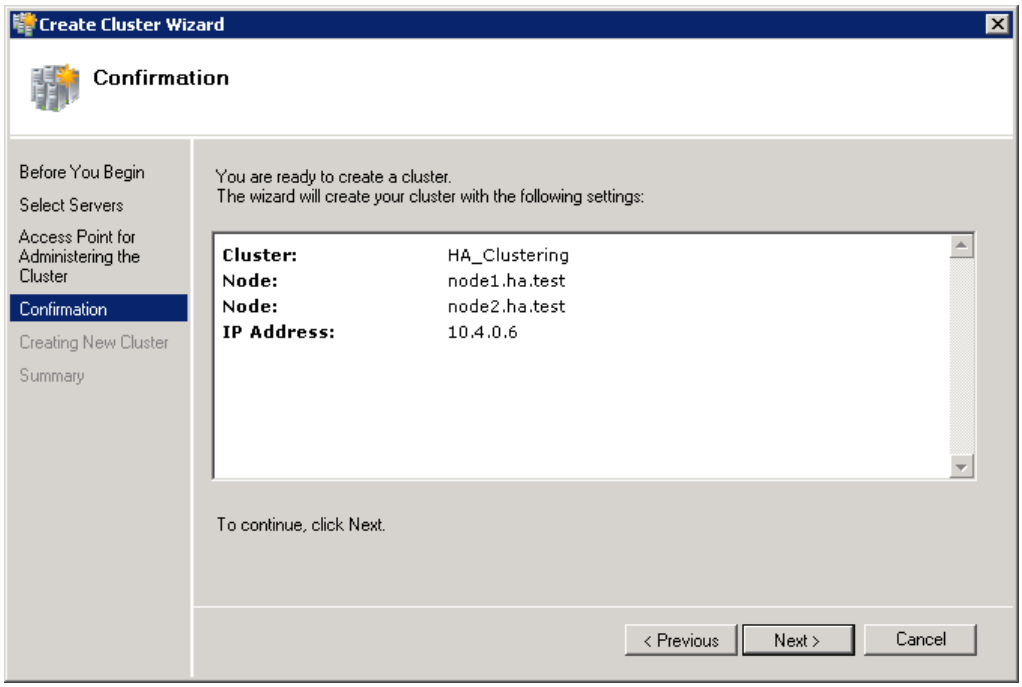

12. The summary is shown after the cluster has created successfully.

After the cluster is created successfully, a new tab with the cluster's name (e.g. HA_Clustering) will be created under [Failover Cluster Management] in the Console tree:

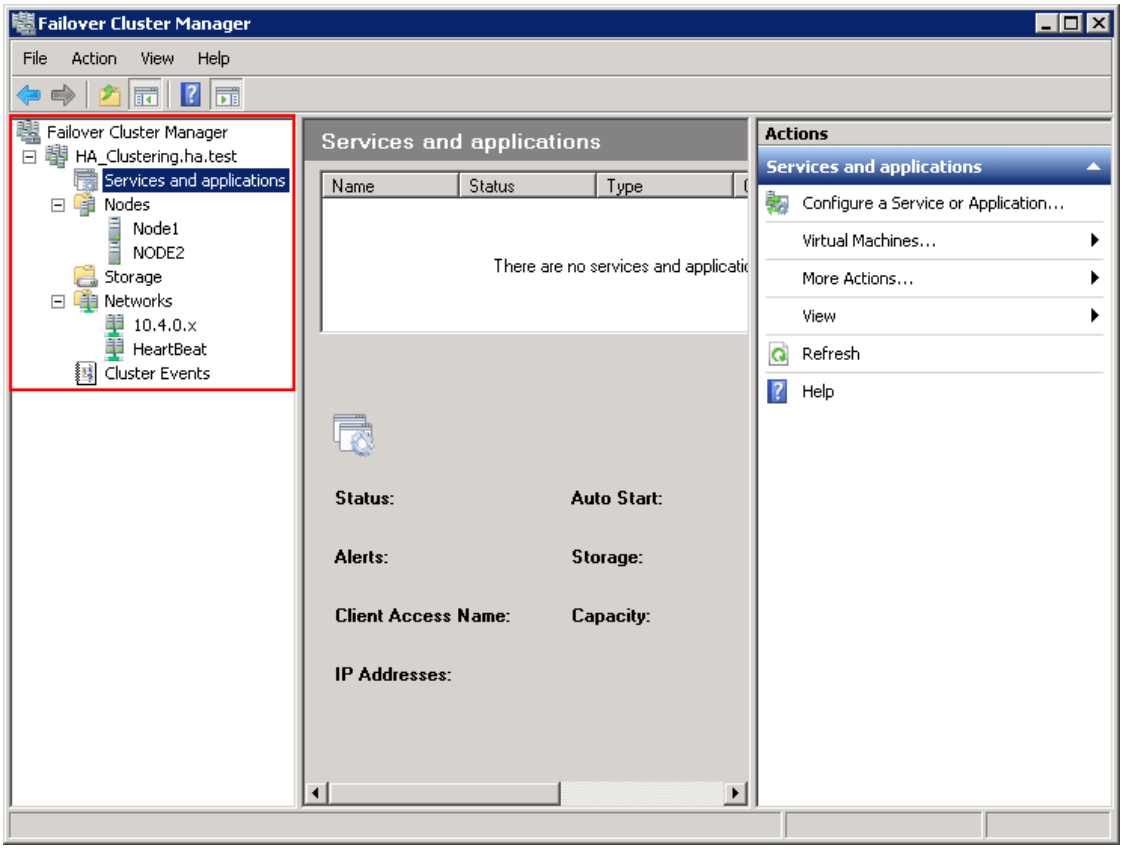

Refer to the following table for more details on functionality of each tab:

Tab	Functions
Services and Applications	This should be empty by default. It holds all services and application that are configured for High Availability in the cluster. E.g. AhsayOBS
Nodes	This provides information such as network on all clusters nodes. In this case, it is Node 1 and Node 2.
Storage	This is a summary of all storages in the cluster. It contains information such as size, owner and status. In our scenario, we should have three: G, H and Q drives.
Network	This summarizes all network connections that relate to the cluster. It include both public and private network configured in the computer.

Step 3 – Configuring Quorum Settings

The quorum determines the number of elements that must be online for the cluster to continue running. Elements are Nodes or disk witness or file share witness. Each element has one vote and the cluster will only continue to operate if it satisfies certain conditions depending on the quorum configuration.

You can configure the quorum setting with the [Failover Cluster Management] interface:

1. Select your cluster, e.g. HA_Clustering.

2. Select [More Action] and click on [Configure Cluster Quorum Setting].

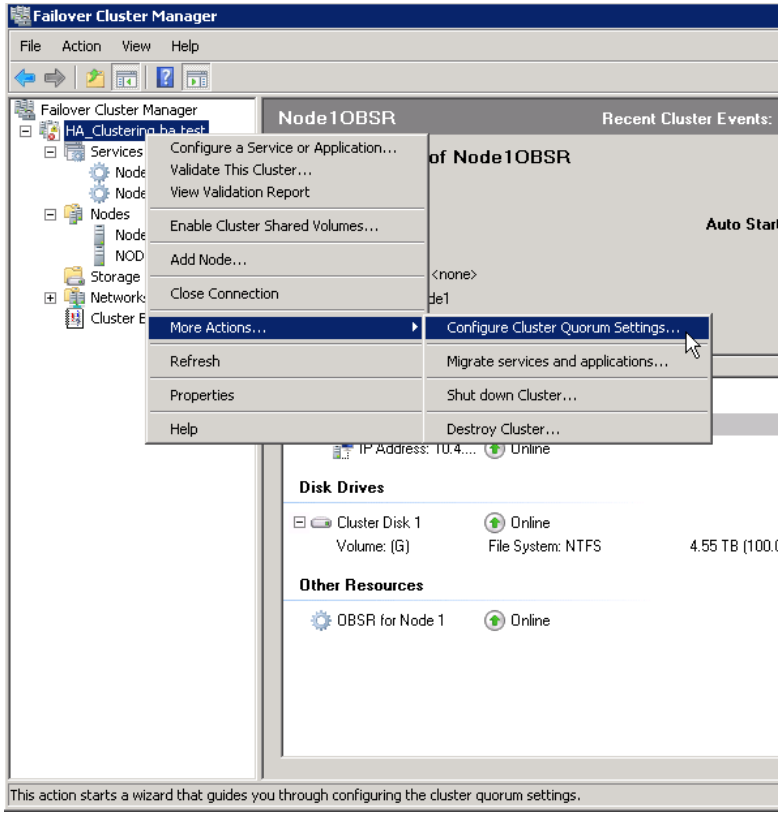

3. Choose [Node and Disk Majority] option for setting up a Two-Node Failover Cluster.

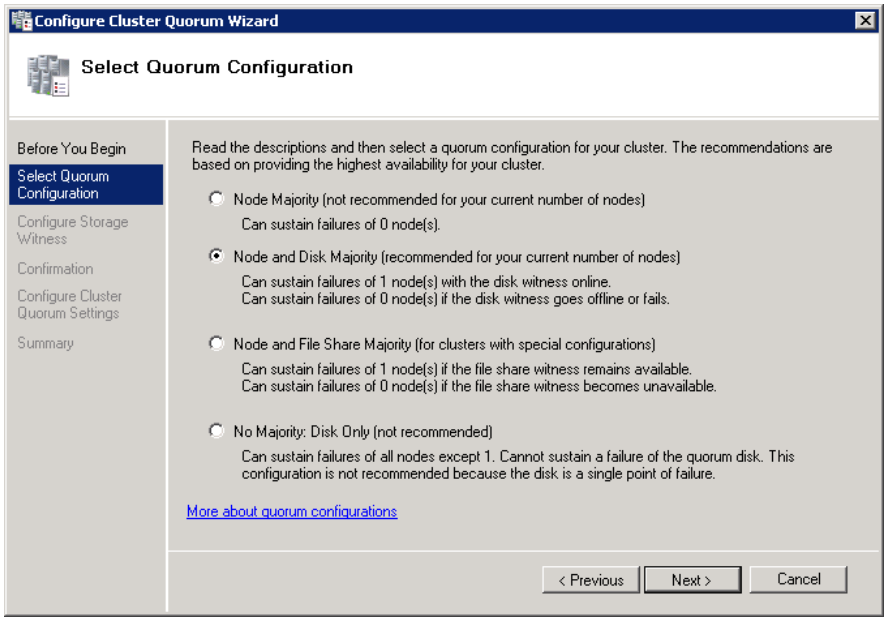

4. Choose Virtual Disk C (e.g. Q Drive) when asked for a storage witness. If you have assigned virtual disk C with a drive letter Q, it should already been assigned to the witness disk role. You can confirm this by viewing the [Storage] section of the cluster.

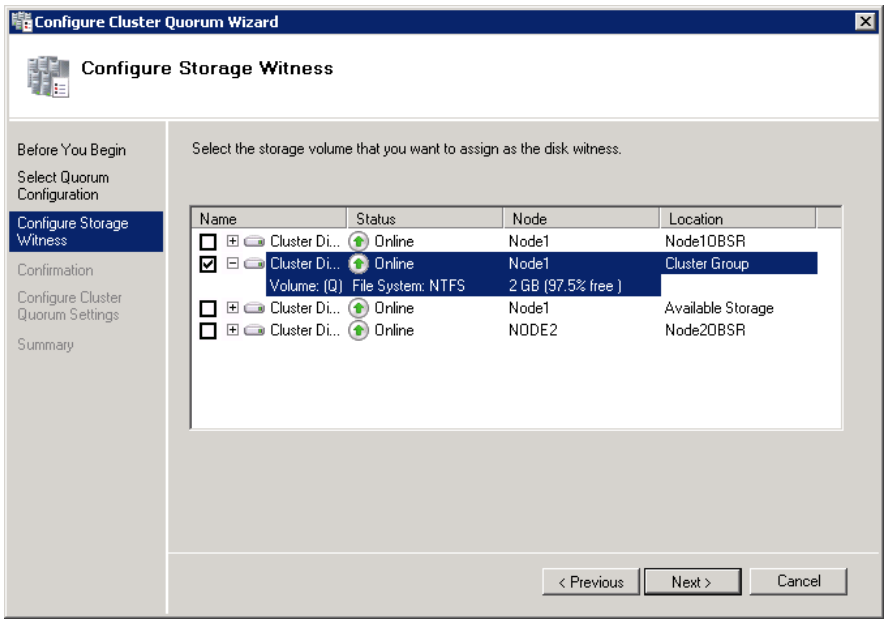

The [Node and Disk Majority] indicates that each Node and witness disk has one vote, so in this scenario, there are 3 votes in total (e.g. 2 Nodes + 1 Quorum).

This means that the cluster will continue to operate if it gets 2 out of 3 votes, in other words, there is a one element fault tolerance before the cluster fails.

For further information on Quorum configuration and settings, please refer to this website.

3.3 Setup first AhsayOBS instance in Windows Platform

To install and setup the first AhsayOBS instance in the cluster, please follow the instructions listed below:

1. Install AhsayOBS using the manual method and skip the step about installing AhsayOBS as a windows service. For further details, please refer to the AhsayOBS and AhsayRPS Setup Guide.

2. Edit startup scripts (startup.bat):

   - Open the startup.bat file under %OBSR_HOME%\bin
   - Under CATALINA_OPTS, change the -Xmx option to a larger value. This option stands for the maximum amount of memory that AhsayOBS can used. E.g. 2G

> For 32-bit Java, the maximum –Xmx settings is between 1.3 - 1.6G. This limitation is removed for 64-bit Java.
>
> If you plan to support a large number of users (greater than 1000) on your system, please consider using the 64-bit Java on 64-bit hardware and increase this value appropriately.

3. Validate server.xml:

   - Open the server.xml file under %OBS_HOME%/conf using a text editor.
   - Under the Connector tag for HTTP and HTTPS, make sure that the port is set properly and address is set to 0.0.0.0.

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<Server port="8014" shutdown="SHUTDOWN">

  <Service name="Tomcat-Standalone">

    <!-- Define a non-SSL HTTP/1.1 Connector on port 8080 -->
    <Connector address="0.0.0.0" port="80"
maxHttpHeaderSize="8192" socketBuffer="8192"
```

```
maxThreads="3000" minSpareThreads="10" maxSpareThreads="50"
maxKeepAliveRequests="100" enableLookups="false"
redirectPort="443" acceptCount="200"
connectionTimeout="60000" disableUploadTimeout="true"
URIEncoding="utf-8"/>

    <!-- Define a SSL HTTP/1.1 Connector on port 8443 -->
    <Connector address="0.0.0.0" port="443"
maxHttpHeaderSize="8192" socketBuffer="8192"
maxThreads="3000" minSpareThreads="10" maxSpareThreads="50"
maxKeepAliveRequests="100" enableLookups="false"
redirectPort="80" acceptCount="200" connectionTimeout="60000"
disableUploadTimeout="true" URIEncoding="utf-8"
scheme="https" secure="true" keystoreFile="conf/keystore"
keystorePass="changeit" clientAuth="false" SSLEnabled="true"
sslProtocol="TLS"/>

    <Engine name="Standalone" defaultHost="localhost" >
      <Host name="localhost" appBase="webapps"
unpackWARs="true" autoDeploy="false" liveDeploy="false">
        <Valve
className="org.apache.catalina.valves.AccessLogValve"
directory="logs" prefix="access_log." suffix=".log"
pattern="common"/>
      </Host>
    </Engine>

  </Service>

</Server>
```

4. Open command prompt and create a new directory for %OBSR_HOME%.

```
>mkdir "G:\AhsayOBS and AhsayRPS"
```

5. Change the working directory to %OBSR_HOME%\bin

```
>cd "G:\AhsayOBS and AhsayRPS\util\bin"
```

6. Run the following command to install the service:

```
>service.exe –i <Service Name> <Service Display Name> <exe
path>
```

Where:

- Parameter *-i* stands for install service and parameter *-r* stands for remove service.
- <Service Name> - is the name of the service.
- <Service Display Name> – is the display name of the service.
- <exe path> - is the path to AobService.exe. By default, it is located at %OBSR_HOME%\bin

For example, in our example, the following command should be run:

```
>service.exe –i Node1OBSR “OBSR for Node1” “G:\AhsayOBS and AhsayRPS\bin\AobService.exe”
```

A message of “Node1OBSR installed” will be displayed on the screen if the service is installed successfully

7. Start AhsayOBS Service by selecting <Service Name> from [Control Panel] > [Administrative Tools] > [Services] and press the [Start] button, e.g. “OBSR for Node1”.

8. Validates AhsayOBS by logging on to it as Administrator using a web browser, e.g. http://10.4.0.4 or http://10.4.0.4:80

Important:
Please do not activate your AhsayOBS license at this stage. Please also restart your computer before you configure AhsayOBS as a generic service.

3.4 Configure AhsayOBS as a generic service

Step 1 – Configure AhsayOBS

In order to use the failover ability of the cluster with AhsayOBS service, you need to create a generic service in the [Failover Cluster Management] Interface.

Please follow the steps below to configure a generic service:

Note:
Assume that AhsayOBS is already installed in Node1.

1. Go to the [Failover Cluster Management] interface at Node 1 by [Start Menu] > [Administrative Tools] > [Failover Cluster Management].

   The ownership of the disk with AhsayOBS installed (G drive) is needed at Node 1, you can check the ownership of storage in the [Storage] tab,

you should see [Volume G] Owned by Node1 in the summary page, as shown below:

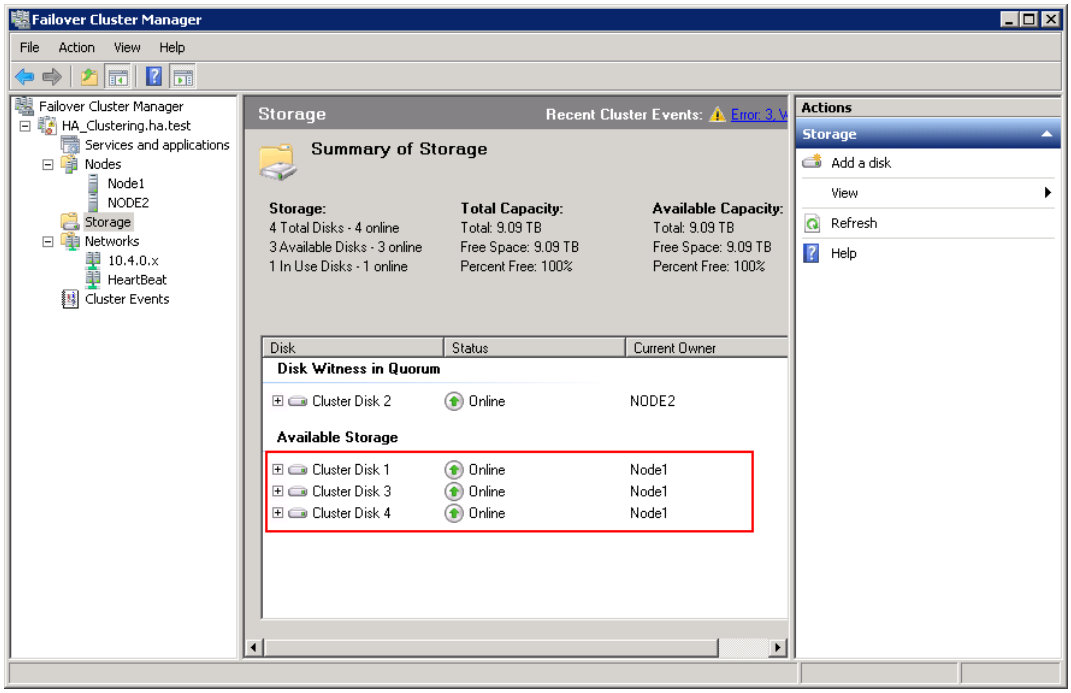

If that is not the case, you can always change the ownership of the disk using the following command:

```
>CLUSTER GROUP "Available Storage" /Move:<Your desire name>
```

For example, to move [Available Storage] to Node 1:

```
>CLUSTER GROUP "Available Storage" /Move:Node1
```

2. 2. Right click on [Service and Application Management] then select **[Configure Service or Application]**.

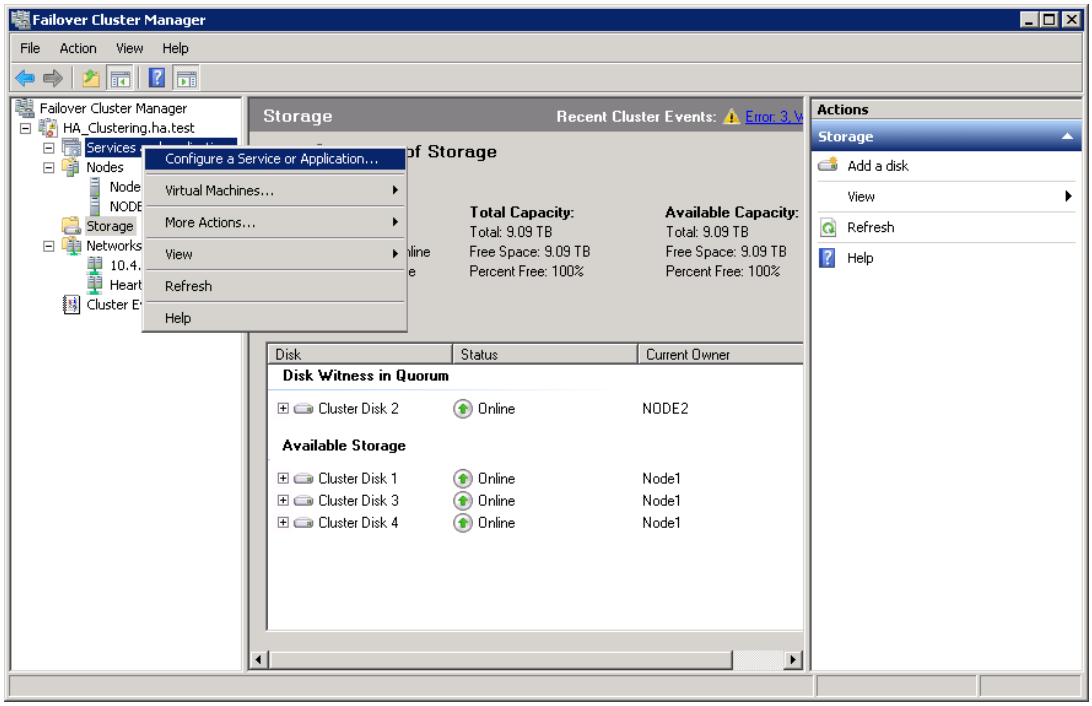

The High Availability Wizard will prompt, below shows steps in configuring AhsayOBS as a generic service.

i. Before you begin.

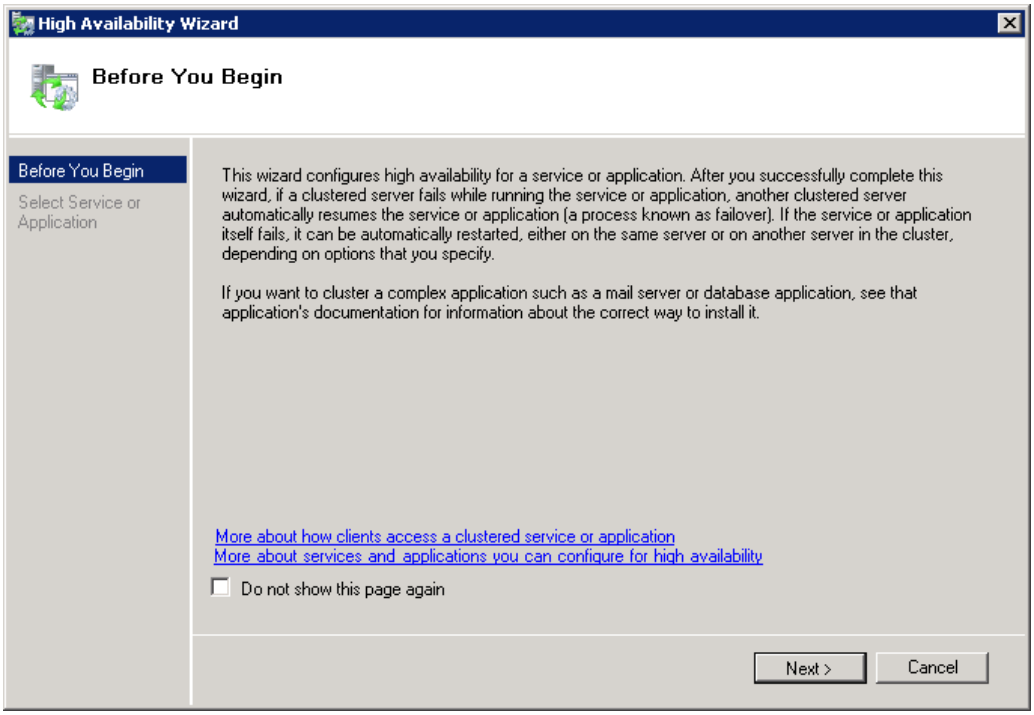

ii. Select [Generic Service].

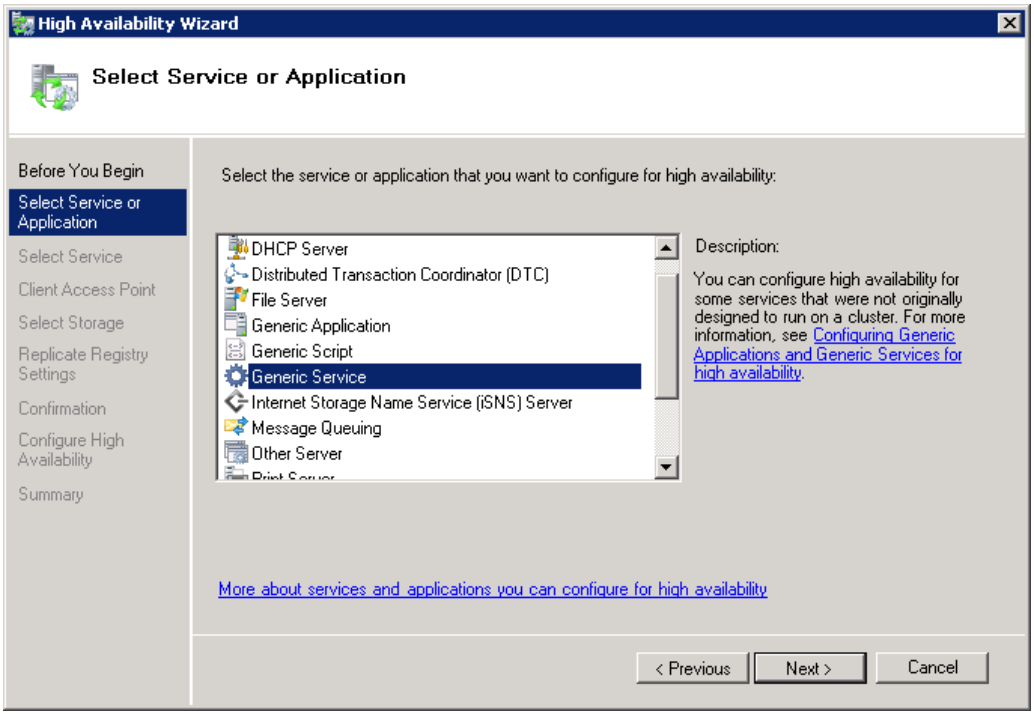

iii. Select AhsayOBS Service from the list, it will be shown with the service's display name (e.g. OBSR for Node1).

If you have renamed it, it will look different from the below screenshot.

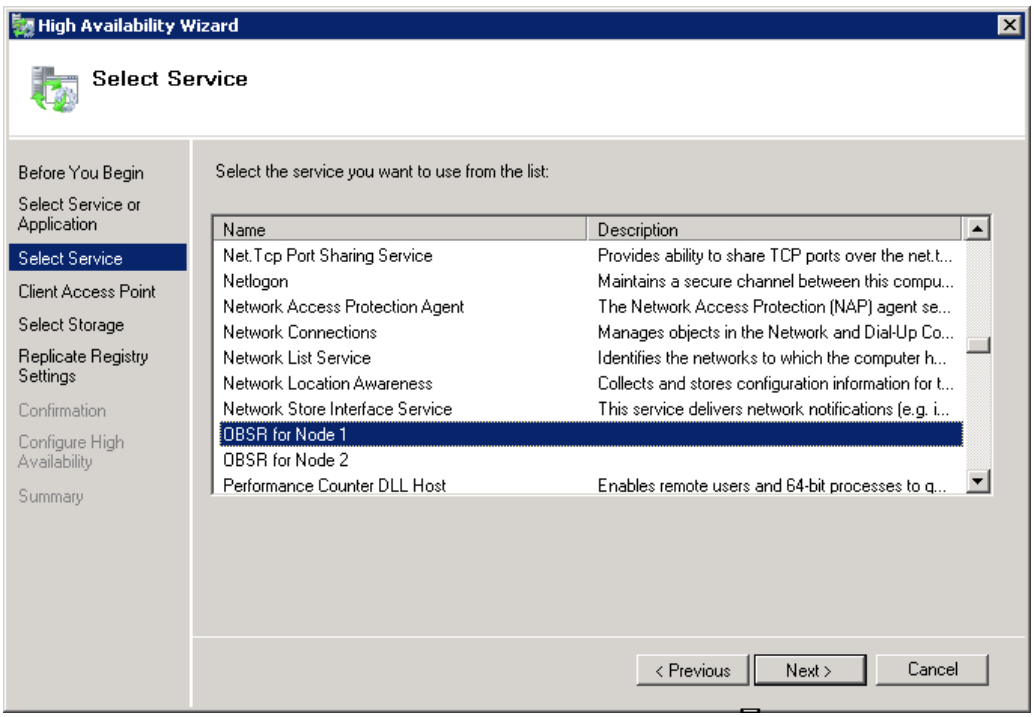

iv. The Name will become this generic service's name on [Failover Cluster Management], therefore it should be unique and informative (e.g. Node1_OBSR).

The Client Access Point is the IP that will be used in accessing the AhsayOBS.

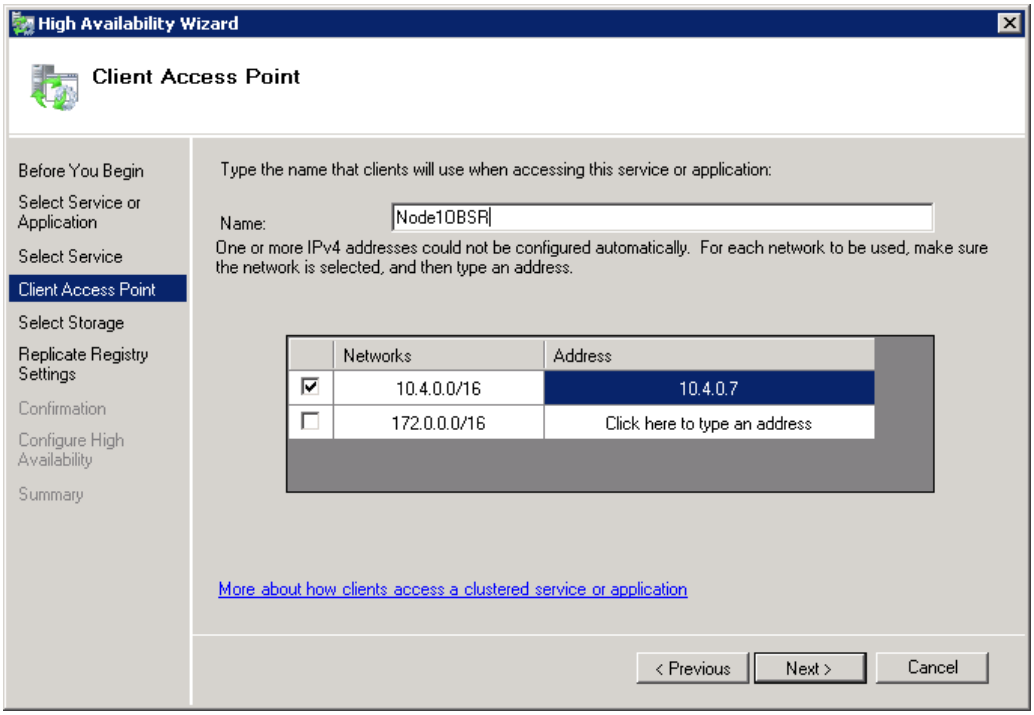

v. Storage select here will follow the generic service, e.g. if the service changes its ownership, all associated storages will also change its ownership.

Choose all storages required to operate AhsayOBS, which includes %USER_HOME%, %SYSTEM_HOME% and %OBSR_HOME%. In this scenario, only Drive G is required.

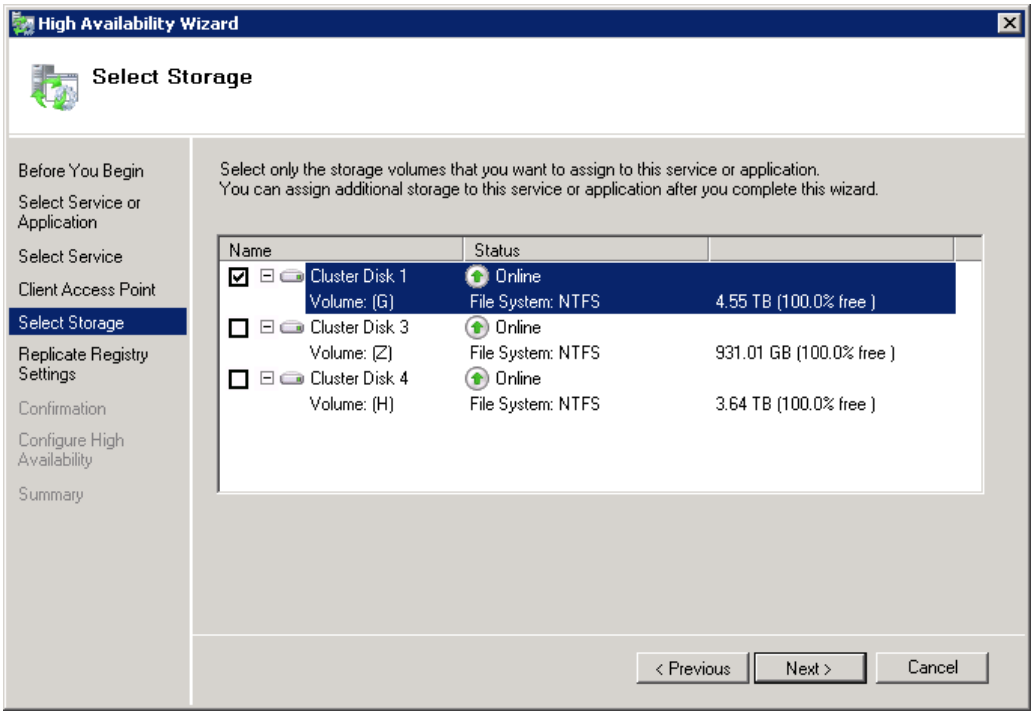

vi. Under the Replicate Registry Settings, some AhsayOBS info will be stored in registry; therefore it is important to have this information available on all nodes.

The location of this key will be dependent on the service name used of the windows service. In general, it is located at:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<Your Service Name>

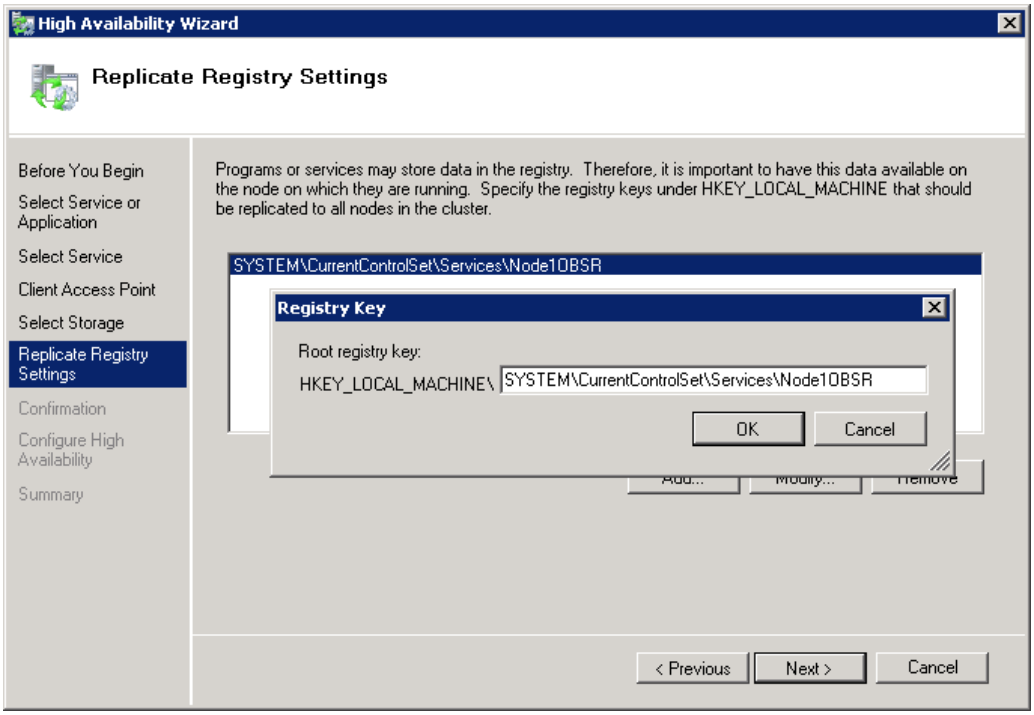

vii. The confirmation screen will be shown after you have finished configuring the registry settings. Click the [Next] button to complete the process.

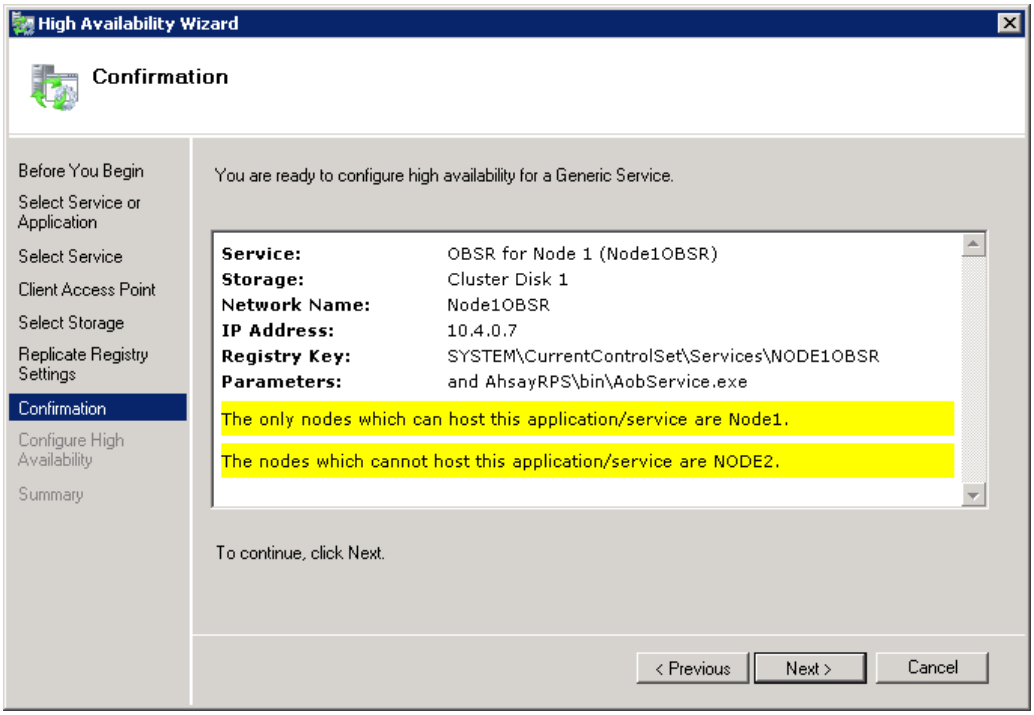

Note:
The summary page may display warning messages suggesting only certain Nodes can access the service. This message can be safely ignored at this stage, as we have not yet setup the other Nodes. Please click [Finish] to complete the process.

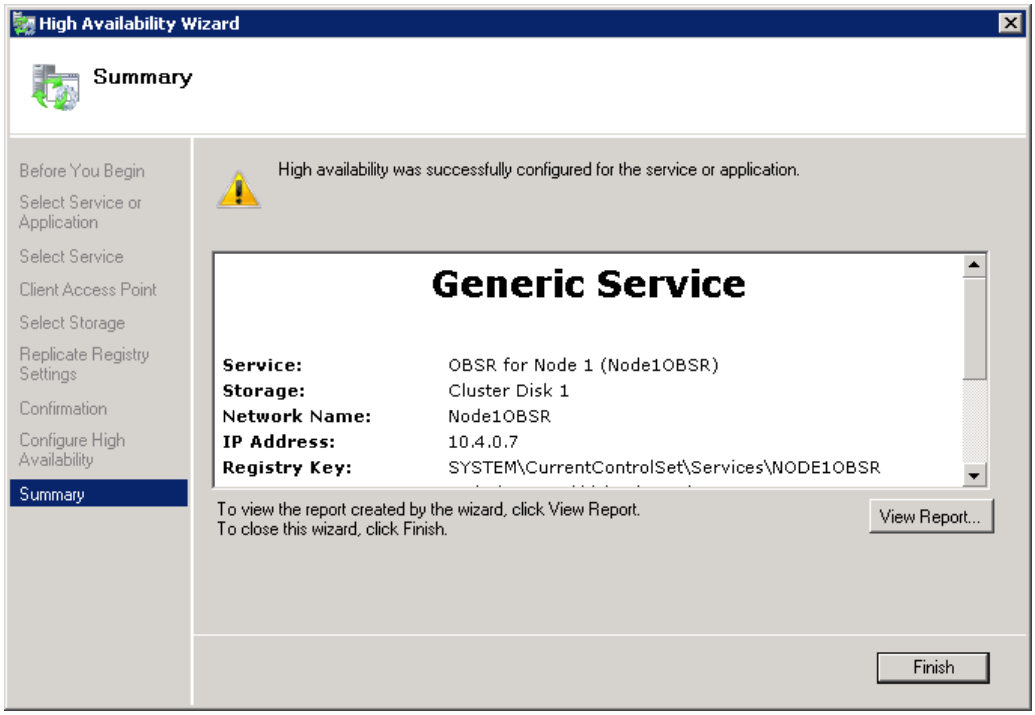

Step 2 – Validate your configuration

- Validate if the generic service is visible in the [Services and Applications] section under [Failover Cluster Management] interface.
- Validate the configuration by browsing the Client Access Point (e.g. http://10.4.0.7:80). This should take you to the AhsayOBS login page.
- (Optional) You can setup a DNS record (e.g. obs1.yourdomain.com) that points to the client access point, e.g. 10.4.0.7.

Important:
You can now activate your AhsayOBS license.

Step 3 – Add / Remove Owner of Cluster resources

Every resource (e.g. Service, IP, storage) has a list of possible owners; a Node can only gain ownership of a resource if it is a possible owner. By default, all nodes are possible owner. This list can be modified in the resource's properties (shown below).

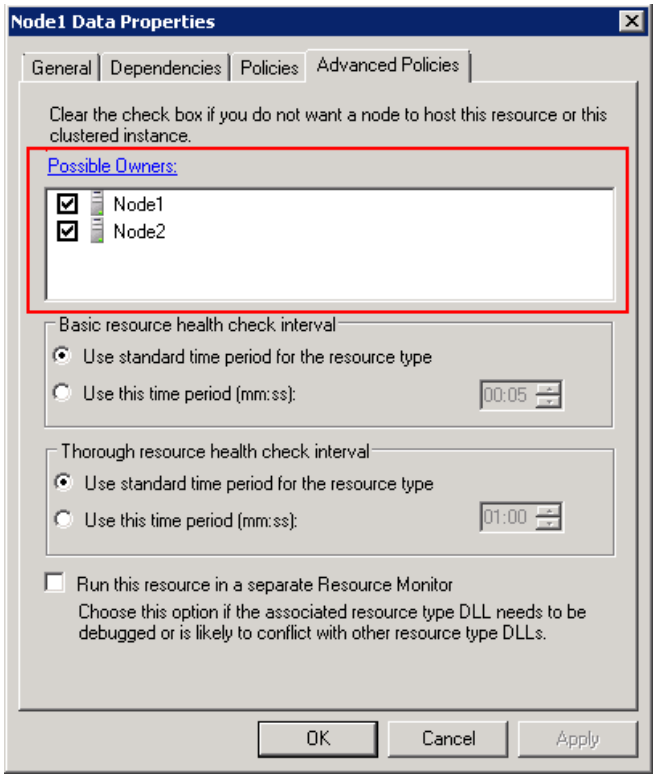

Step 4 - Registering the generic service to other Node

To register a generic service to a Node:

1. Trigger failover once to the other node (e.g. from Node 1 to 2):

   Right click on the generic service > [Move this service or application to another node].

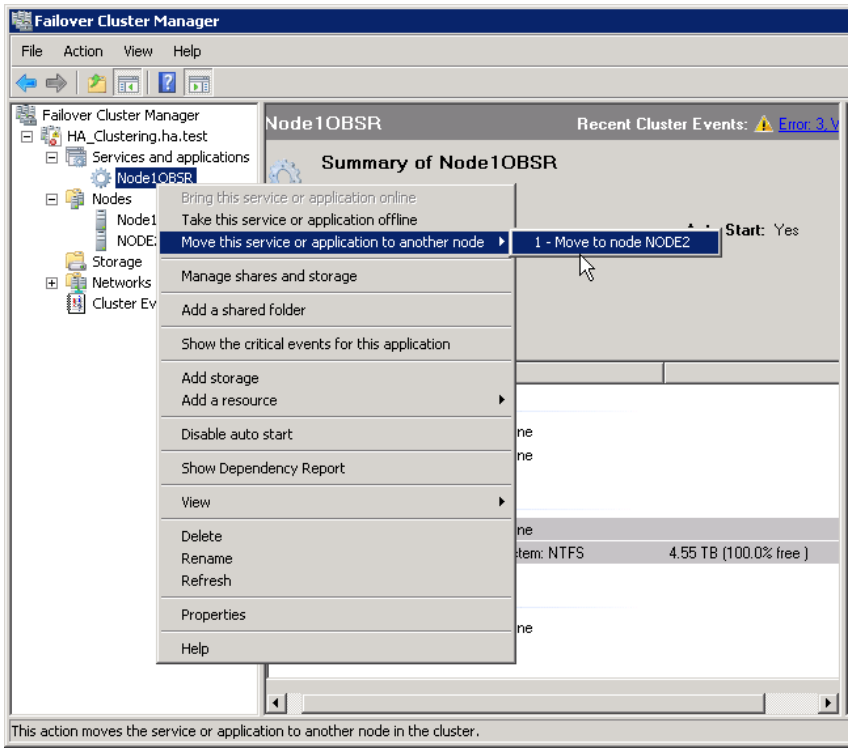

   The failover is expected to fail as service cannot be started on the Node 2, but this will add the required registry information on Node 2.

2. Restart Node 2, this will add the service on Node 2. You can validate this by observing the list of service after Node 2 has been restarted, and the new service should have been installed into the list.

3. Validate by triggering failover to the Node again, this time the Node is registered, the failover should succeed.

3.5 Setup additional AhsayOBS in the cluster

The procedure in setting up additional AhsayOBS is exactly the same as setting up the first instance, but it must satisfy the following conditions:

- The port numbers must be unique from all other AhsayOBS in the cluster.
- The service name must be unique from all other AhsayOBS in the cluster.
- The server port must be unique from all other AhsayOBS in the cluster.
- Each disk should only be associated with one instance of AhsayOBS, as a virtual drive can be owned by only one Node at any given time.

The steps for setting up additional AhsayOBS in the cluster are summarized as below, where Node1_OBSR has already been done throughout the guide.

Generic service Name	Client Access Point	Service Name	Installation Location	Port Used
Node1_OBSR	10.4.0.7	Node1OBSR	G:\AhsayOBS and AhsayRPS	80, 443 & 8014
Node2_OBSR	10.4.0.8	Node2OBSR	H:\AhsayOBS and AhsayRPS	70, 8443 & 8015

Please perform the following steps on Node 2:

1. Install AhsayOBS in H:\AhsayOBS and AhsayRPS
2. Rename the service name to a unique, non-default name. E.g. Node2OBSR
3. Validate if it can access by go to the Administrator Login Page via web browser. Do not update License at this stage.
4. Setup Node2OBSR as a generic service named Node1_OBSR and Client access point being 10.4.0.8.
5. Validate Node2_OBSR by invoking http://10.4.0.8:70, you can now activate its license.
6. Register Node2_OBSR on Node 1.
7. Both Node1OBSR and Node2OBSR are now running high availability in the failover cluster.

4 Verify AhsayOBS HA Setup on Windows

After finishing setting up the High Availability Option of AhsayOBS, you must verify the following sections before putting the servers into a production environment.

4.1 Cluster Setup

To verify the cluster configuration of Windows 2008, please connect to the AhsayOBS instances by using its cluster IP address (or DNS name). After you have entered the Client Access Point, e.g. http://<cluster-ip-address>, in a web browser, you should be able to view the Administrator Login page.

4.2 Cluster Failover

There are two methods to verify the failover feature of the cluster. However, it is recommended to perform both tests to ensure the cluster is working properly.

Method 1 - Manual Failover

A failover event can be simulated manually with the aid of [Failover Cluster Management] on windows. Please do the following to trigger a manual failover:

1. Open [Failover Cluster Management].

2. Right Click on the <Generic Service> (e.g. Node1OBSR) and Click on [Move this service or application to another node].

3. Click on the <Generic Service> and you should be able to see the Current Owner is switched to Node2 in the Summary page.

4. Invoke the AhsayOBS Web Admin Console, you are expected to see the standard Administrator page again.

If generic service, Node1OBSR was running at Node1, you can failover the server to Node2 with Failover Cluster Management. During the failover, the disk with OBSR installed will detach from Node 1 and reattach to the other Node 2, which will start the Node1OBSR service.

Method 2 – Auto Failover

To verify the failover ability of the cluster, we should also simulate the event of a server breakdown. Please follow the instructions below:

1. Logon to the Node 1.
2. Shutdown the Node 1, go to [Start] > [Shutdown].
3. Logon to the Node 2.
4. Open [Failover Cluster Management].
5. Click on the <Generic Service> (e.g. Node1OBSR), the Current Owner should be switched to Node2 in the summary page.
6. Invoke the AhsayOBS Web Admin Console, you are expected to see the standard Administrator page again.

4.3 Other Verification Tests

After you have finished the tests mentioned above, you can verify the AhsayOBS cluster using the AhsayOBM/AhsayACB as you would normally.

- Check that AhsayOBM and AhsayACB can backup and restore.
- Check web restore applet.
- Check whether your branding is correct.
- Check that the email reports for admin and users are correct.
- Check whether data integrity could remain when backup is interrupted by server restarts and temporarily network disconnection.

5 Setup High Availability AhsayOBS on Linux

This chapter describes how to setup the AhsayOBS high availability on Linux Platform.

AhsayOBS High Availability setup is summarized as follows:

1. Operating system and storage configuration
2. Failover cluster configuration.
3. Setup the first AhsayOBS instance.
4. Configure AhsayOBS as a high availability service.
5. Setup additional AhsayOBS instance(s) to the cluster. (This step can be repeated for as many AhsayOBS Instances for the cluster as required.)

5.1 Operating system and storage configuration

When the R710 servers and MD3000 are connected, install Red Hat Enterprise Linux 5 AP onto each of the cluster node. During the installation process, please choose to install the [Clustering] features, which include most of the essential components for running high availability. In addition, please uncheck the [Virtualization] feature during installation, as that might prevents Dell's drivers from installing.

It is highly recommended to disable SELinux, it is known that the cluster software does not work with SELinux Enforce mode/Permissive mode for RedHat Enterprise Linux 5.3 AP.

After that, please install the SAS 5/E drivers and multi-path driver on each node, these are available in the MD3000 Resource disk.

Below shows steps in installing SAS 5/E drivers and multi-path driver on Linux environment.

1. Ensure that the packages libXp, libXtst, gcc and kernel development (kernel-devel) are installed, as they are required to complete the installation.

   Run the following commands to ensure that the packages are installed:

   ```
   [root] # yum install libXp libXtst gcc kernel-devel
   ```

   Note:
   The version of kernel-devel MUST match your running kernel. For example, if your running kernel is 2.6.18-128.el5 then kernel-devel-2.6.18-128.el5 is required. All available packages can be found on Red Hat's Website.

2. Run the linux installation script in the MD3000 resource disk with the commands:

   ```
   [root] # mkdir /mount/cdrom
   [root] # mount /dev/cdrom /mount/cdrom
   [root] # sh /mount/cdrom/linux/install.sh
   ```

   Below menu should be displayed:

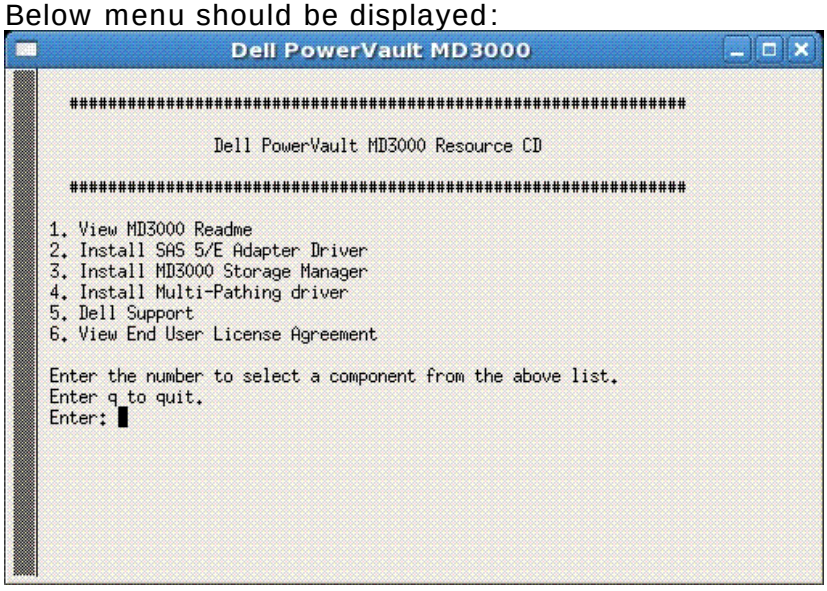

3. Select the [2. Install SAS 5/E Adapter Driver] option, this must be installed prior installing the multi-path driver. The following message appears:

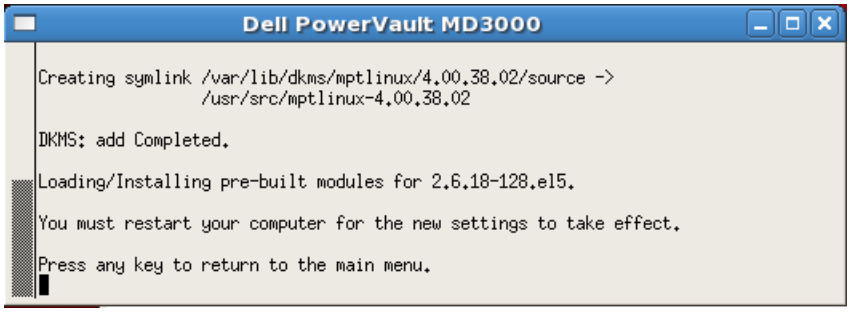

Please restart the computer before continuing to the next steps.

4. Select [4. Install Multi-Pathing driver] option, if the current running kernel version is not supported, please try downloading the latest Resource disk from Dell's website.

Successful will display

```
Creating new MPP initrd image...
        You must now edit your boot loader configuration file, /boot/grub/menu.l
t, to
        add a new boot menu, which uses mpp-2.6.18-128.el5.img as the initrd ima
e.
        Now Reboot the system for MPP to take effect.
        The new boot menu entry should look something like this (note that it ma

        vary with different system configuration):

        ...

                title Red Hat Linux (2.6.18-128.el5) with MPP support
                root (hd0,5)
                kernel /vmlinuz-2.6.18-128.el5 ro root=LABEL=RH9
                initrd /mpp-2.6.18-128.el5.img
        ...
MPP driver package has been successfully installed on your system.
make: Leaving directory `/usr/src/linuxrdac-09.03.0C06.0234'

/etc/modprobe.conf: added alias reference for 'mppUpper'
depmod....

Saving old initrd as /boot/initrd-2.6.18-128.el5_old.img
Making new initrd as /boot/initrd-2.6.18-128.el5.img
(If next boot fails, revert to the _old initrd image)
mkinitrd..........................................................

DKMS: install Completed.

Your DKMS tree now includes:
linuxrdac, 09.03.0C06.0234, 2.6.18-164.el5, i686: built
linuxrdac, 09.03.0C06.0234, 2.6.18-164.el5, x86_64: built
linuxrdac, 09.03.0C06.0234, 2.6.18-128.el5, i686: built
linuxrdac, 09.03.0C06.0234, 2.6.18-128.el5, x86_64: installed
linuxrdac, 09.03.0C06.0234, 2.6.18-92.el5, i686: built
linuxrdac, 09.03.0C06.0234, 2.6.18-92.el5, x86_64: built

Is this computer going to be used as an
 Oracle (TM) Real Application Cluster (RAC) node (y/n)?
```

5. The following steps instructs the kernel to boot with the multi-path driver:

   i. Confirm the ram disk file name.

```
[root] # ls /boot/ | grep mpp
mpp-2.6.18.el5.img
```

   ii. Edit grub.conf file.

```
[root] # vi /etc/grub.conf
```

iii. Edit it as shown below:

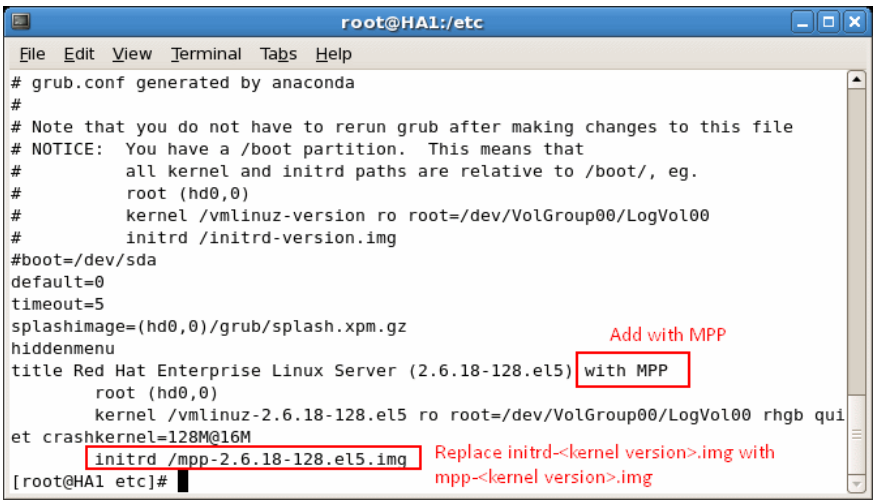

iv. Reboot the computer

Note:
Please ensure that kernel 2.6.18-128 is used for Red Hat Enterprise Linux 5.3. The Multi-Pathing driver for MD3000 works with this kernel only. To check the kernel version, please use:

```
uname -r
```

If the kernel is updated to another version by the auto-update feature, please update the /etc/grub.conf file and put the entries related to 2.6.18-128 kernel in front of other version and reboot the machine. This will force the machine to boot up with 2.6.18-128 kernel automatically.

Building your virtual disks

Install Dell's Modular Disk Storage Manager onto one of the cluster node, this is used for MD3000 disk array configuration.

The Dell PowerVault MD3000 Disk Array System comprised 15 individual disks, though independently they can be integrated into a virtual disk and be used as is a single disk. As shown in the diagram in the chapter 1.6, three virtual disks are created, two of which will be used for installing AhsayOBS, and one will be used as a quorum disk, which will be covered later in this document.

Please follow the instructions below to build the virtual disk:

1. Open [Modular Disk Storage Manager].

2. Go to the [Configure] tab and select [Create Disk Groups and Virtual Disks] option.

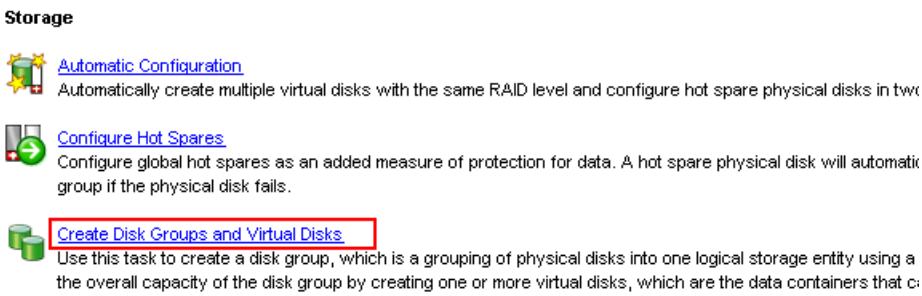

3. Choose [Disk group] and click [Next].

4. Choose the desire RAID level and number of disk.

5. After that, partition the disk group into virtual disks.

The following settings are used in the Active/Active Two-Node Failover Cluster in our example:

Virtual Disk	Disk Group	RAID Level	Size	Host Group
Virtual Disk A	Disk 0 – 2	RAID5	1.8TB	Host Group 1
Virtual Disk B	Disk 3 – 5	RAID5	1.8TB	Host Group 1
Virtual Disk C	Disk 6 – 7	RAID1	2GB	Host Group 1

Note:
Not all available space has to be used.

Mapping virtual disk to Host Group
Each cluster Node connected to MD3000 is seen as a Host, in order for a Host to access a virtual disk, we need to explicitly map the disk to the Host.

1. Create host group by selecting [Create Host Group] option under [Configure] tab, enter a name (e.g. Host Group 1), add both cluster Node 1 and Node 2 into the Host Group and then click [OK].

2. Select [Create Host-to-Virtual Disk Mappings] option under [Configure] tab, select [Host Group 1] and map Virtual Disk A, B and C to it.

Now all Virtual Disks are mapped to Hosts in Host Group 1, so that both cluster nodes are able to access these virtual disks.

5.2 Failover cluster configuration

This section provides guidance to setup the cluster using Conga.

Step 1 – Preparing cluster nodes for Conga

Before doing this step, please make sure that the [RHEL Clustering] and [RHEL Cluster-Storage] module are enabled in the corresponding system profile on RedHat Network Portal.

At each cluster node, update ricci, ricci is the client agent of **Conga**, it broadcast node information to luci, installed at the management Node.

```
[root] # yum update ricci
```

Configure ricci to start on boot:

```
[root] # chkconfig ricci on
```

Start ricci service:

```
[root] # service ricci start
```

Execute the following commands on the management node to install Conga server luci:

```
[root] # yum update luci
```

Note:
The luci can be configured on any node (including the cluster node).
In production environment, it is highly recommended to install luci on a dedicated management node that is not part of the cluster. However, for testing purpose, it is setup on one of the cluster node in the example.

Initialize the luci server and assign the admin password:

```
[root] # luci_admin init
```

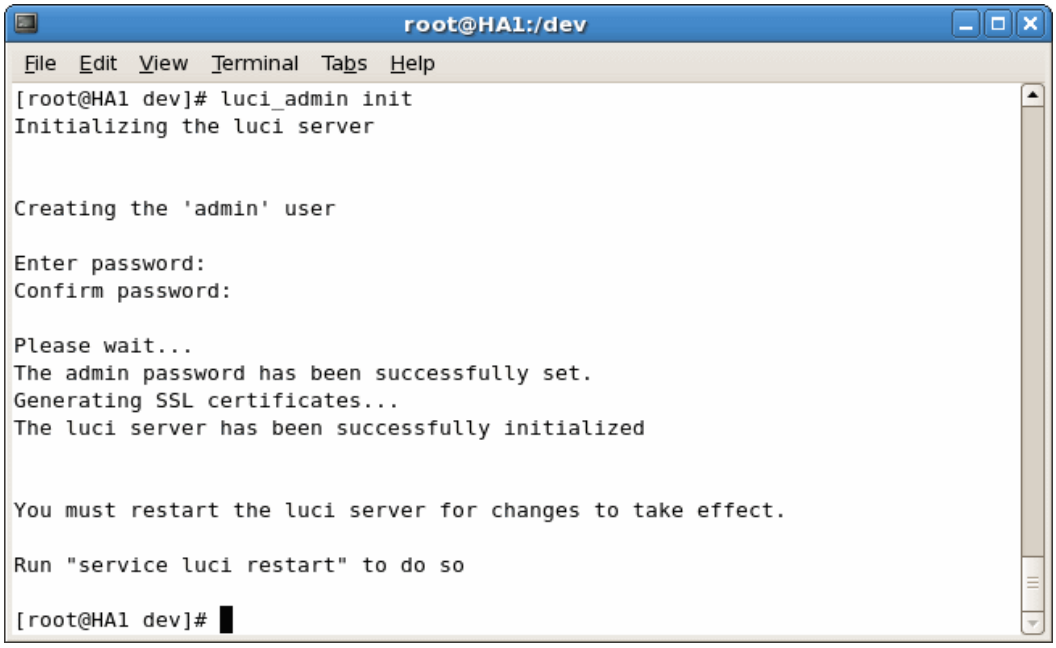

Configure luci to start on boot:

```
[root] # chkconfig luci on
```

Start the luci service:

```
[root] # service luci start
Starting luci: Generating https SSL certificates… done [ OK ]

Point your web browser to https://management.example.com:8084
to access luci
```

Please use this url to connect to luci server:
https://<management_node_hostname_or_IP>:8084

Turn on cluster locking on both nodes:

```
[root] # lvmconf --enable-cluster
```

Step 2 – Creating the cluster with Conga

1. Connect to luci server with any browser. The following page should be shown:

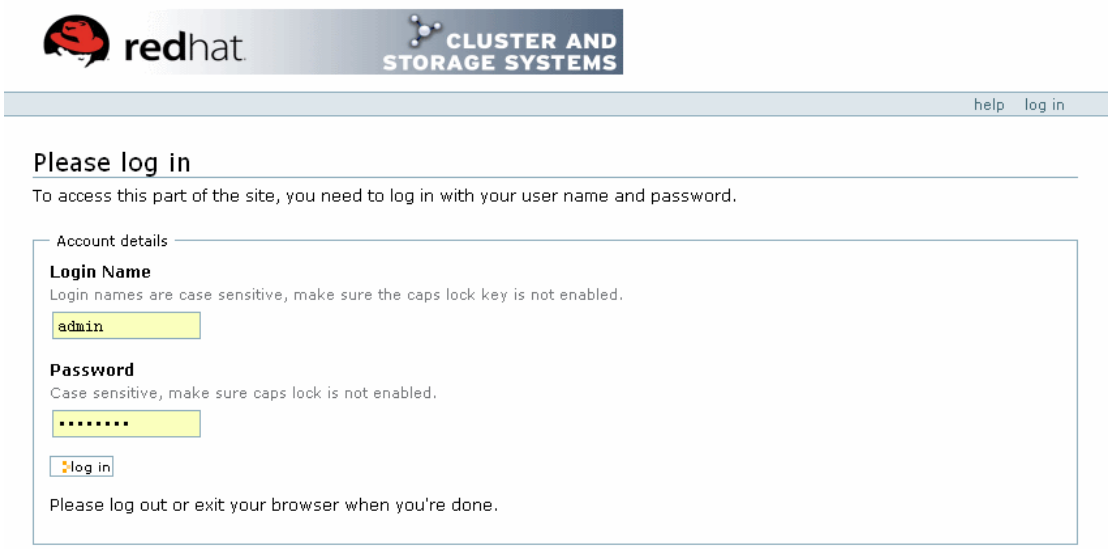

2. Login to luci with admin, click on the [Cluster] tab. Then click [Create a New Cluster], enter the cluster information. E.g. Cluster name and node information.

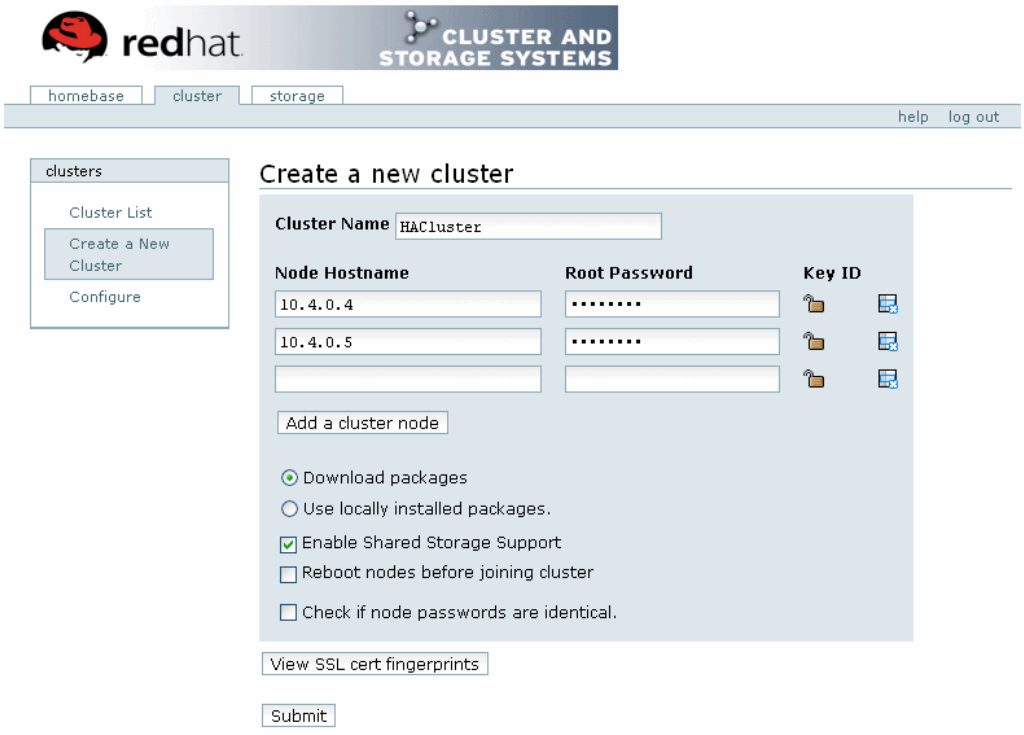

3. After the cluster is created, you will be redirected to the [General] tab of the cluster.

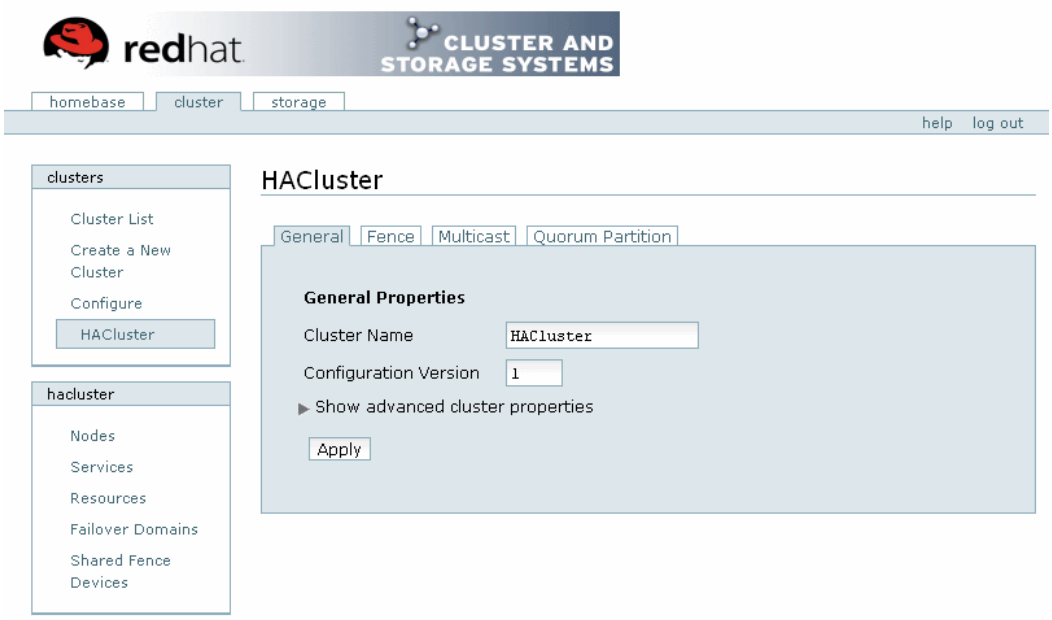

4. Start running the cluster manager (CMAN) by using:

```
[root] # service cman start
```

Step 3 – Configuring cluster storage with Conga

Please ensure that the multi-path drivers are installed before continuing. The followings steps are required to perform on one of the cluster nodes only. This is because the all the nodes have access to the same virtual disk.

1. Click on the Storage tab on luci, the cluster nodes should be listed under "System List", shown as below.

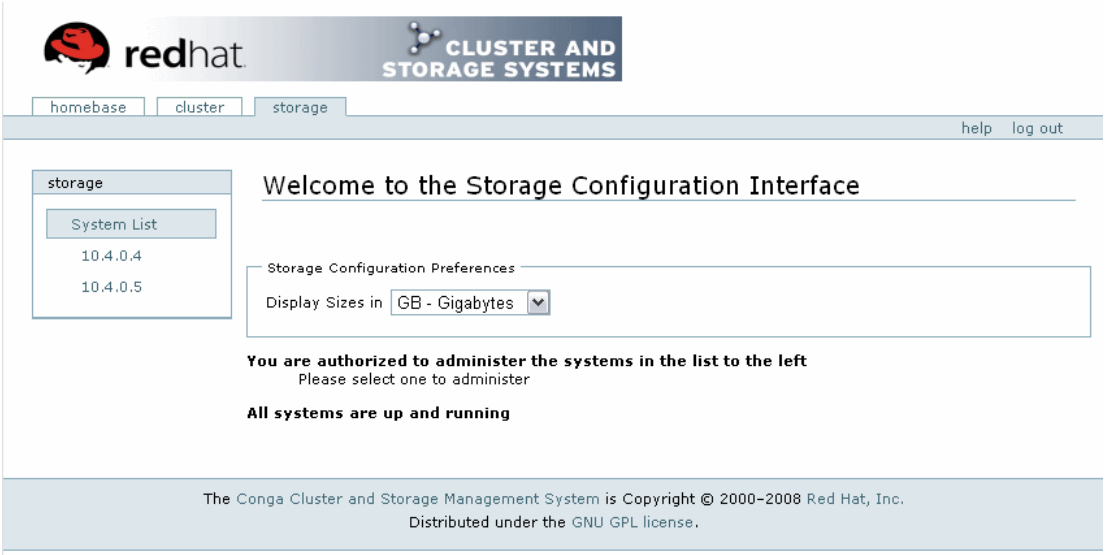

2. In the "System List" section, select a Node. The hard drives that are visible to that Node will be displayed. Check the configuration on all nodes and ensure that they have the same view of storage. If this is not the case, please click the [Reprobe Storage] button. If this does not fix the issue, it is advised to revise the steps in Operating System and storage Configuration.

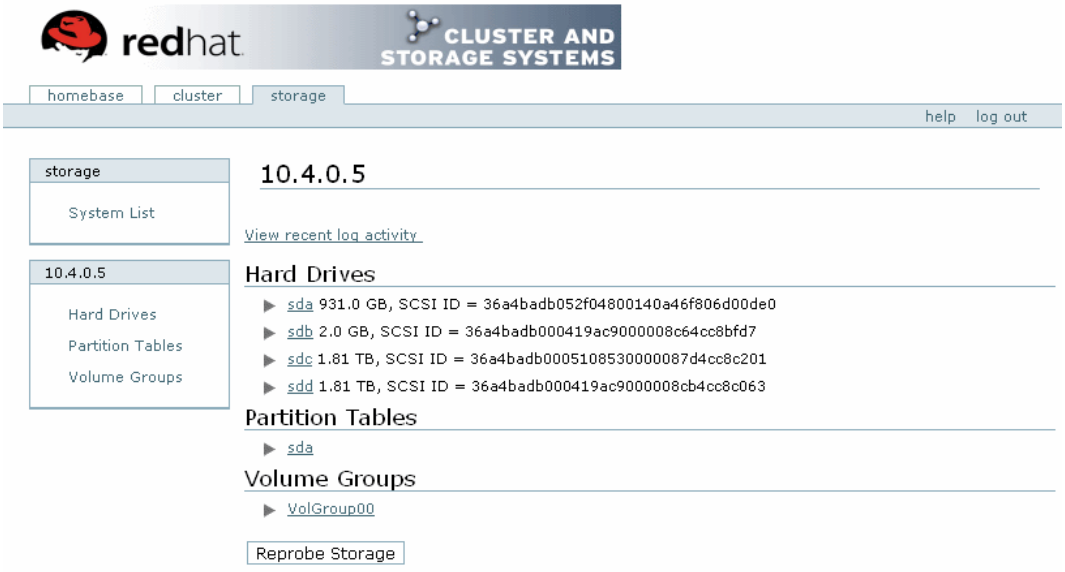

3. To partition the shared storage, select a node from “System List”, then click on [Partition Tables]. Select [New Partition Table].

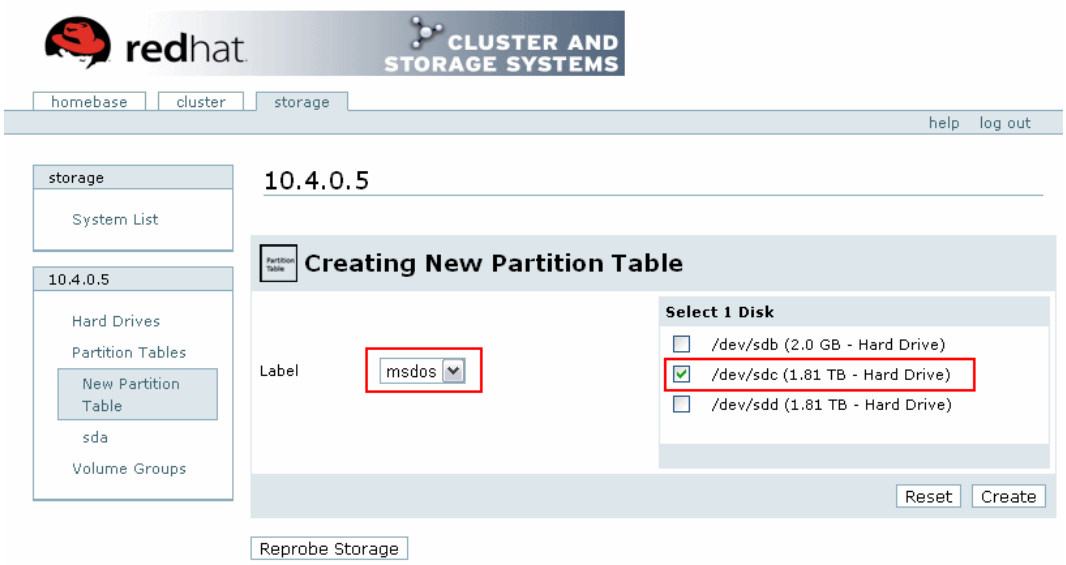

In the Label field, select “msdos” option, then check the box of your virtual disk and click the [Create] button.

4. Click on the [Unused Space], leave the [Content] field as Empty and then click the [Create] button.

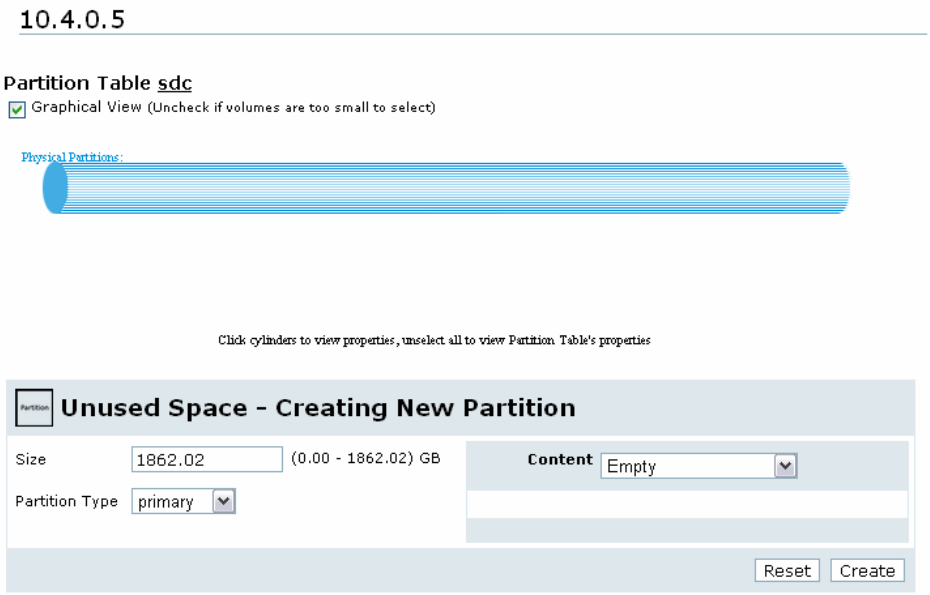

5. On the left panel, click [Volume Groups] and click [New Volume Group]. Then enter a name for the volume group. (E.g. *VG1*) Ensure that the [Clustered] option is set to “True”, and select the newly created data partition at step 4 above. After that, click the [Create] button.

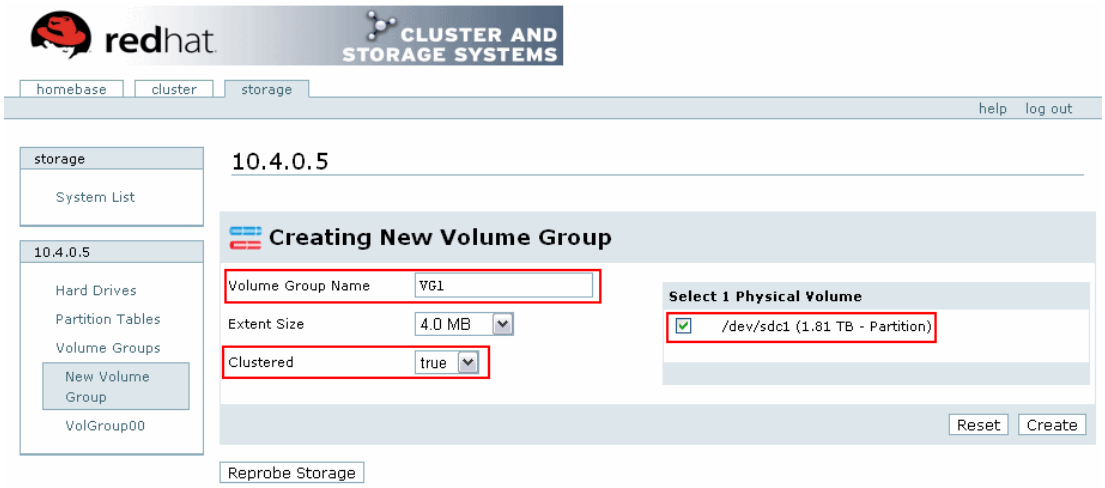

6. After Volume group has been created, click on [New Logical Volume] located at bottom. Enter a Logical Volume Name (E.g. *LV1*), choose "Linux Extended FS" for the [Content] field, please leave [List in /etc/fstab] option as "false", and click the [Create] button.

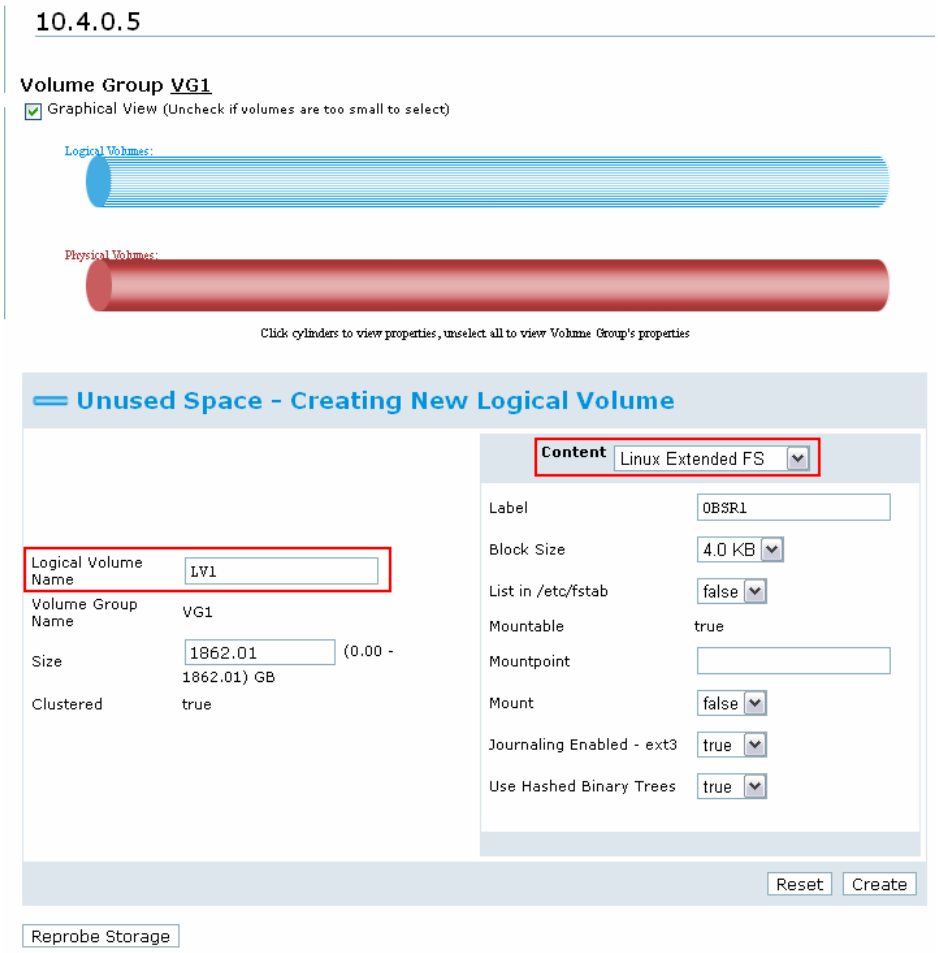

Note:
It is required to restart the cluster lvm daemon on all nodes. Run the following command to restart the daemon:

```
[root] # service clvmd restart
```

7. Repeat the above steps for all volumes you wish to create, except for the quorum partition which will be cover in the next section. Below shows the resultants view after all volumes described in the scenario Section are created.

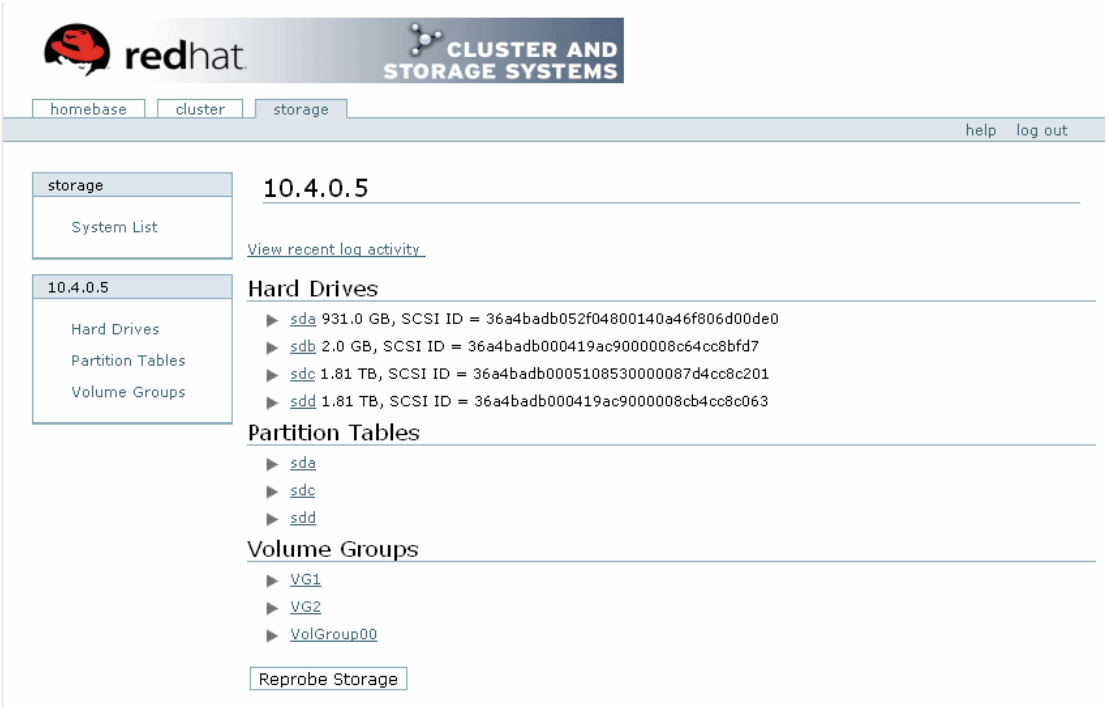

Note:
If error occurs when trying to create a volume group or logical volume. Please go to the [Storage] > [<Node>] and verify if the storage settings are the same. If not, please do one of the followings:

- Go to [storage] > [<node>] and click the [Reprobe Storage] button on all nodes.
- SSH to the nodes and use the following command:

```
[root]# partprobe
[root]# service clvmd stop
[root]# service clvmd start
```

Step 4 – Configuring the quorum partition

Every Node has its own Cluster Manager (CMAN) that keeps track of cluster quorum by monitoring the count of cluster nodes. If more than half the nodes are active (n/2 + 1), the cluster has quorum. However in a two node cluster this algorithm will still results in 2, meaning that at least 2 nodes must be running in order for the cluster to be establish a quorum, which is against the initiative for high availability.

Therefore for a 2 node cluster, a quorum partition is required. A quorum partition is a small disk partition shared across the cluster; it adds an extra vote towards the cluster resulting in a total vote of 3. In that case, the cluster can establish a quorum when ever 2 out of 3 votes are acquired. Hence even if one node is down, the cluster still remains.

A small shared partition is required to create a quorum. In this example, a 2GB virtual disk, /dev/sdb, is used. To create a quorum using Conga, please follow the steps below:
Note:
A raw partition of 10MB is the recommended size

1. In the [storage] Tab of luci, select a node from “System List”, then click on [Partition Tables]. Select [New Partition Table].

In the [Label] field, select “msdos” option, then check the box of your virtual disk and click the [Create] button.

2. Click on the [Unused Space] tab, leave the [Content] field as Empty and then click the [Create] button.

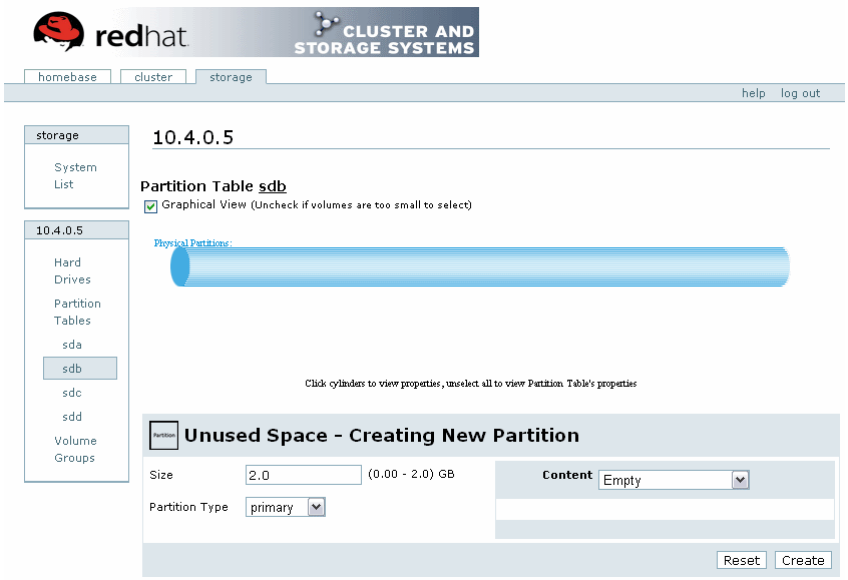

3. Use the *mkqdisk* command to create the quorum partition, this is only need to run on one of the cluster nodes.

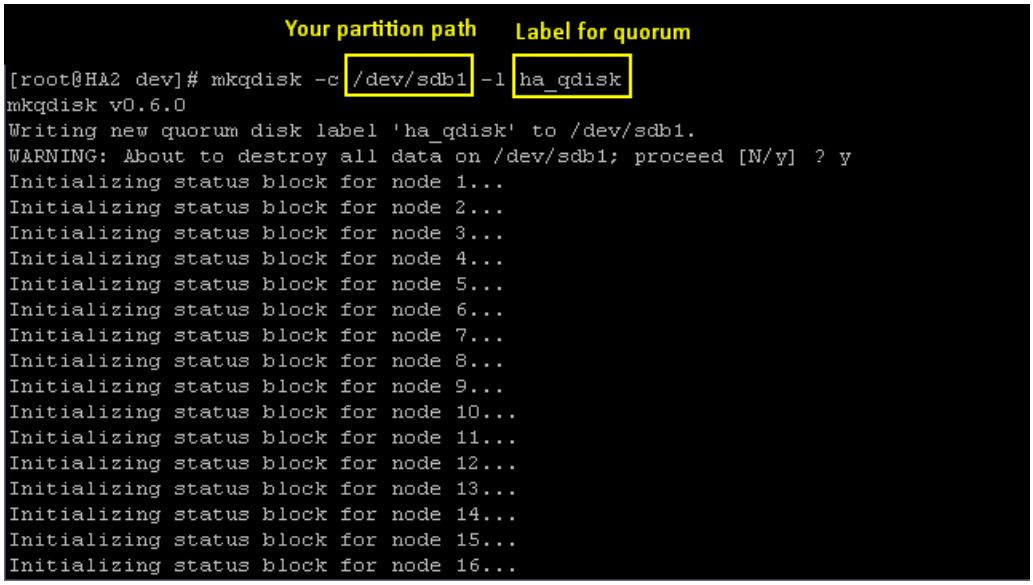

4. To confirm that the quorum partition has been created, please run this command on both nodes:

```
[root]# mkqdisk –L
```

If the quorum partition is not visible on a node, please run the following command:

```
[root]# partprobe
```

5. Configure the quorum partition using conga.

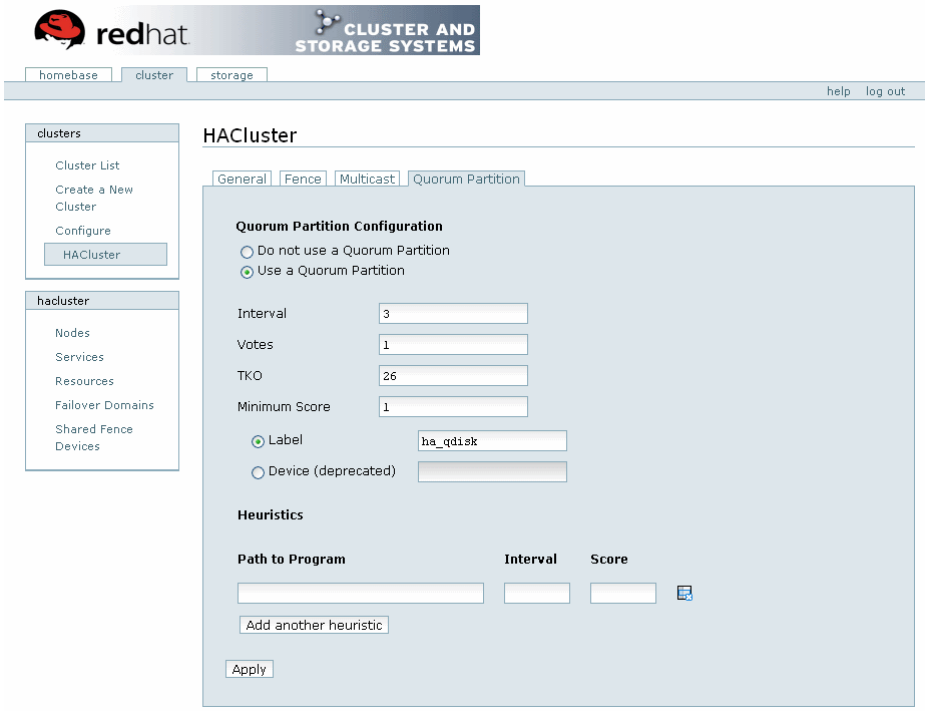

The values of [Interval] and [TKO] were chosen to allow time for the multi-path failover, for more information on configuring the quorum partition, please read the manual of qdisk:

```
[root] # man qdisk
```

Start quorum disk daemon with command:

```
[root] # service qdiskd start
```

Go back to the [Cluster] page, the [Total Cluster Votes] and [minimum Required Quorum] should be updated:

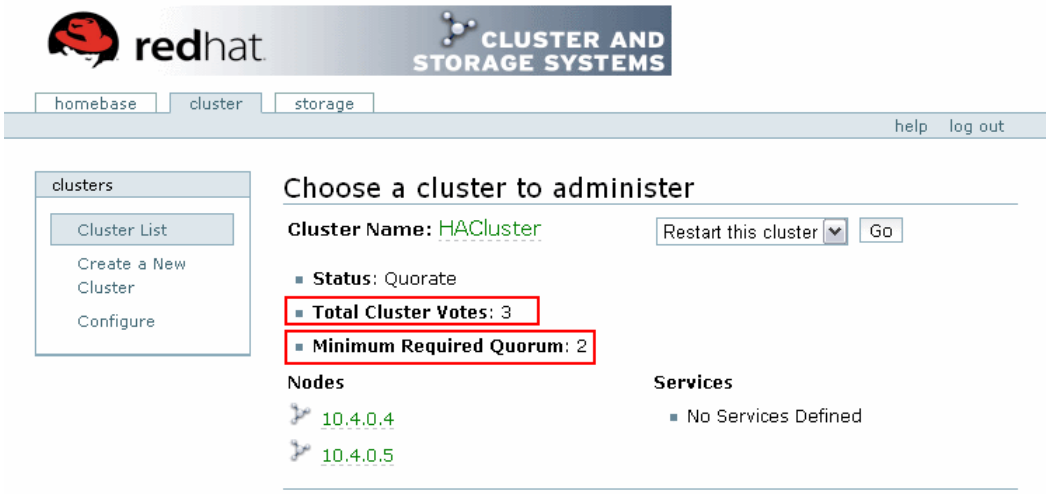

6. Configuring Heuristic setting:

In a two-node cluster, a special occasion needs to be considered. When the network between the two nodes failed, both nodes will lose communication with each other but at the same time stayed connected with the quorum disk, hence CMAN in both nodes believes the other node has failed and tried to fence the other node from the cluster.

To avoid this situation, a heuristic can be added to evaluates a node's network connectively and remove itself form the cluster if its network fails, the typical heuristic would be to ping the network router.

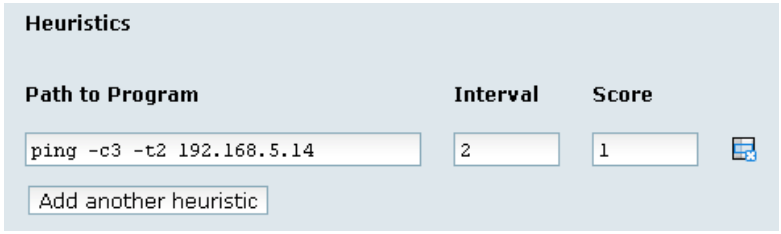

Step 5 – General cluster configuration

To ensure smooth failover, some of the default parameters need to be modified. We need to define a quorum_dev_poll that is larger than the qdisk timeout defined in the previous step (Interval * TKO). The quorum_dev_poll is the maximum time (in milliseconds) that CMAN would allow qdisk to not communicate before declaring it to be dead.

On one of the cluster node:

1. Go to the */etc/cluster/* directory and modify cluster.conf.

2. Add the following to the cman tab:

```
<cman quorum_dev_poll="100000">
```

3. Increment the config_version under the [cluster] tab by 1

4. Save changes, then deploy the updated cluster.conf with the command:

```
[root] # ccs_tool update /etc/cluster/cluster.conf
```

In addition, the maximum time CMAN allows a node to be under votes before losing quorum should be defined to be greater than the quorum_dev_poll, this can be done with conga:

5. Select the cluster in luci, on the [General] tab; click the arrows beside "Show advanced cluster properties". Configure it as shown below and click the [Apply] button.

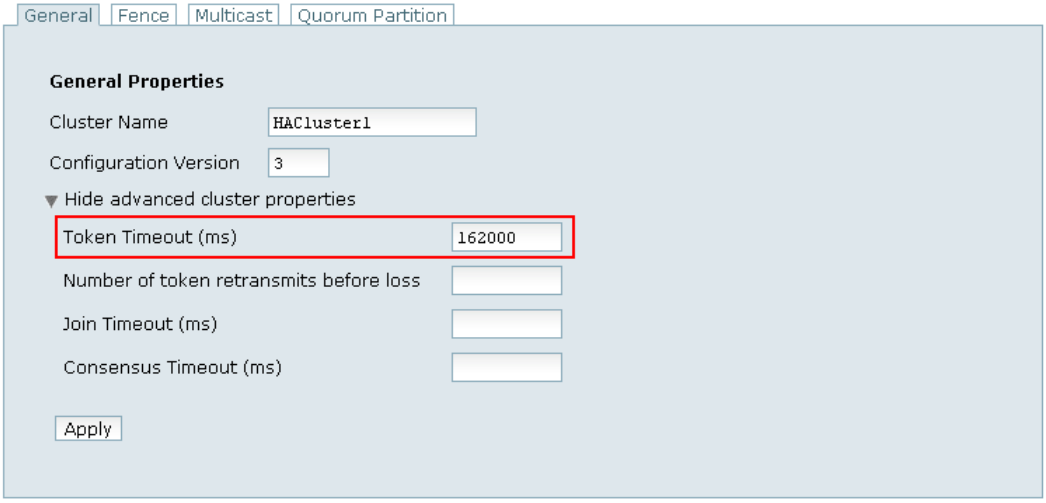

This will give the CMAN enough buffering time in the event of failover. For more information on cluster parameter setting, please refer to the link below:

https://access.redhat.com/kb/docs/DOC-2882

6. Configure all necessary daemon to start on boot

```
[root] # chkconfig qdiskd on
[root] # chkconfig clvmd on
[root] # chkconfig cman on
[root] # chkconfig rgmanager on
```

Step 6 – Setup fencing device (Recommended)

Fencing is the disconnection of a node from the cluster's shared storage. Fencing cuts off I/O from the shared storage, thus ensuring data integrity.

The Red Hat Cluster Suite provides a variety of fencing methods, for this example, Dell's iDRAC6 enterprise is used as the fencing device, which is capable of powering off inoperable nodes.

Please repeat these steps on every node to add a fencing device:

1. Go to the "Nodes" pages under cluster tab, select a Node.

2. Under Main Fencing method, click on [Add a fence device to this level].

3. Select the fencing device from the list and enter appropriate information.

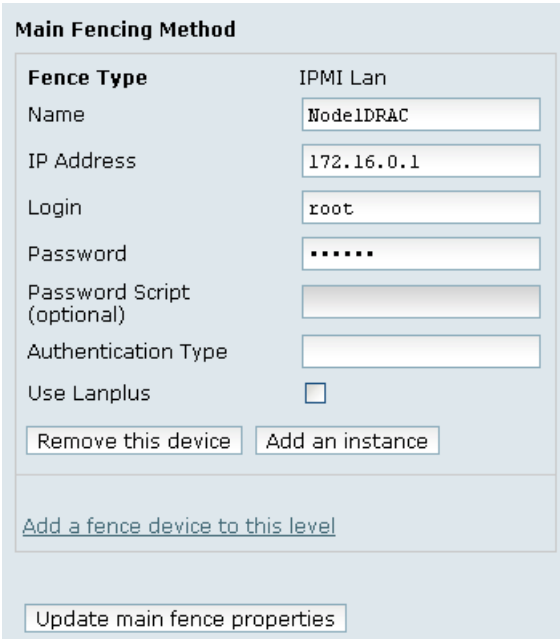

4. Click Update main fence properties to complete the setup

To test the fencing mechanism, invoke:

```
[root] # fence_node {fully qualified hostname or ip address
of <Node> }
```

If successful, the Node is expected to reboot.

5.3 Setup first AhsayOBS instance in Linux Platform

To install and setup the first AhsayOBS instance in the cluster, please follow the instructions listed below:

1. If there is an existing version of AhsayOBS running, stop AhsayOBS and rename the %OBSR_HOME% to %OBSR_HOME_BAK%.

2. On one of the cluster node (e.g. Node1), mount one of the partition to a specify mount point. In this example, */dev/VG1/LV1*, will be mounted to */mnt/mount1.*

3. Install/Upgrade AhsayOBS. For further details, please refer to the AhsayOBS and AhsayRPS Setup Guide.

4. Stop AhsayOBS service with the command:

```
[root] # service obsr stop
```

5. Disable obsr service for automatic startup with the command:

```
[root] # chkconfig obsr off
```

6. Rename the script */etc/init.d/obsr* with a non-default unique service name. e.g. */etc/init.d/obsr1*

7. Modify the renamed script as specify in Appendix B.

8. Deploy the script to the */etc/init.d* directory of the other node, ensure the permission of the script is 755.

9. Edit startup shell scripts (startup.sh):

   i. Open the startup.sh file under %OBSR_HOME%/bin

   ii. Under CATALINA_PID, change "/var/run/obsr.pid" to "/var/run/< Service name>.pid". E.g. /var/run/obsr1.pid

   iii. Under CATALINA_OPTS, change the "-Xmx" option to a larger value. This option stands for the maximum amount of memory that AhsayOBS can used. E.g. 2G

   For 32-bit java, the maximum –Xmx settings is between 1.3-1.6G. This limitation is removed for 64-bit java. If you plan to support a lot of users (greater than 1000) on your system, please consider using the 64-bit java on 64-bit hardware and increase this value appropriately. Assume that each user requires 1MB plus a 50% total buffer.

   If you receive any "Out of Memory" error for Java, you need to increase this setting.

10. Edit shutdown shell scripts (shutdown.sh):

    i. Open the shutdown.sh file under %OBSR_HOME%/bin

    ii. Under CATALINA_PID, change "/var/run/obsr.pid" to "/var/run/< Service name>.pid". E.g. /var/run/obsr1.pid

11. Validate server.xml:

    i. Open the server.xml file under %OBS_HOME%/conf using a text editor.

    ii. Under the Connector tag for HTTP and HTTPS, make sure that the "port" is set properly and "address" is set to "0.0.0.0".

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<Server port="8014" shutdown="SHUTDOWN">

  <Service name="Tomcat-Standalone">

    <!-- Define a non-SSL HTTP/1.1 Connector on port
8080 -->
    <Connector address="0.0.0.0" port="80"
maxHttpHeaderSize="8192" socketBuffer="8192"
maxThreads="3000" minSpareThreads="10"
maxSpareThreads="50" maxKeepAliveRequests="100"
enableLookups="false" redirectPort="443"
acceptCount="200" connectionTimeout="60000"
disableUploadTimeout="true" URIEncoding="utf-8"/>

    <!-- Define a SSL HTTP/1.1 Connector on port 8443 --
>
    <Connector address="0.0.0.0" port="443"
maxHttpHeaderSize="8192" socketBuffer="8192"
maxThreads="3000" minSpareThreads="10"
maxSpareThreads="50" maxKeepAliveRequests="100"
enableLookups="false" redirectPort="80"
acceptCount="200" connectionTimeout="60000"
disableUploadTimeout="true" URIEncoding="utf-8"
scheme="https" secure="true"
keystoreFile="conf/keystore" keystorePass="changeit"
clientAuth="false" SSLEnabled="true" sslProtocol="TLS"/>

    <Engine name="Standalone" defaultHost="localhost" >
      <Host name="localhost" appBase="webapps"
unpackWARs="true" autoDeploy="false" liveDeploy="false">
        <Valve
className="org.apache.catalina.valves.AccessLogValve"
directory="logs" prefix="access_log." suffix=".log"
pattern="common"/>
      </Host>
    </Engine>
```

```
  </Service>

</Server>
```

12. On the Node where the installation was done, invoke the following command to start AhsayOBS service:

```
[root] # service <Service Name> start
```

13. Validates AhsayOBS by logging on to it as Administrator using a web browser, e.g. http://10.4.0.4 or http://10.4.0.4:80.

Important:
Please do not activate your AhsayOBS license at this stage.

5.4 Configure AhsayOBS as a high availability service

Step1 – Configure AhsayOBS

In order to use the failover ability of the cluster with AhsayOBS service, it can be done by creating a high availability service with Red Hat Cluster Suite with the conga interface.

Please follow the steps below to configure a high availability service:

Note:
Assume that AhsayOBS is already installed with Node1 and the device /dev/VG1/LV1 is mounted on Node1

1. Stop AhsayOBS service with the command:

```
[root] # service <Service Name> stop
```

2. Login to luci as administrator.

3. Click on the [cluster] tab and select the cluster, under the cluster's name at the left hand panel, select [Services].

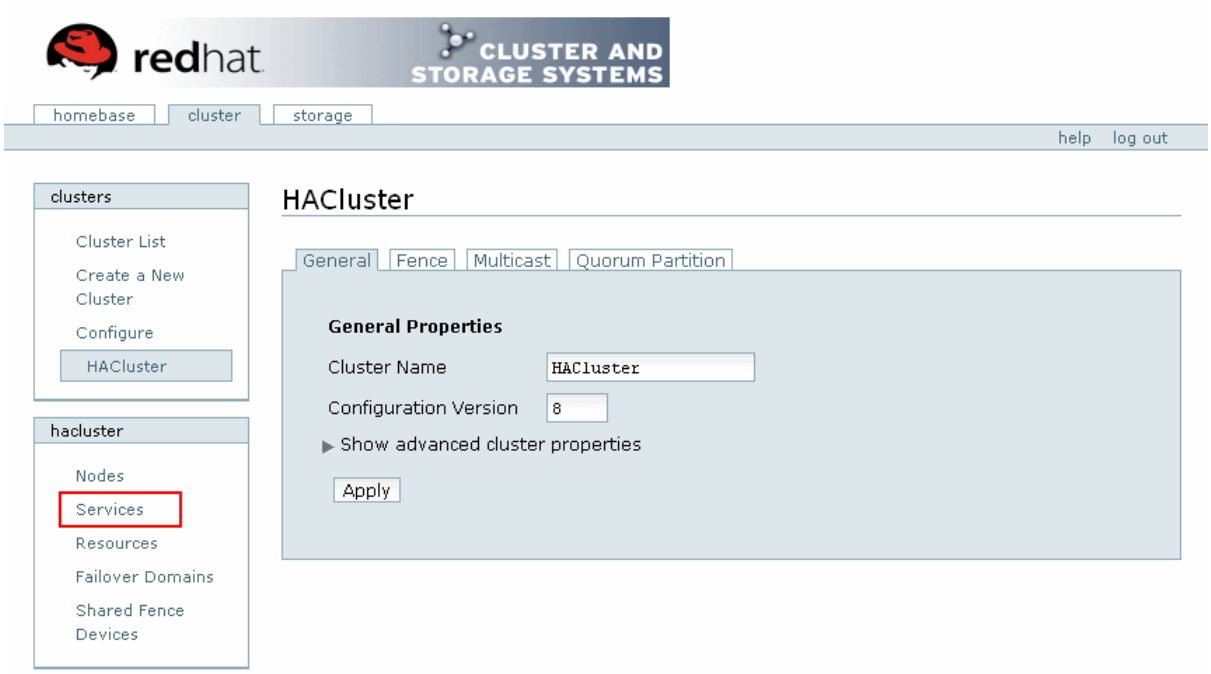

4. From the expanded option under [Services], select [Add a service], give it an appropriate <Service Name> and select “relocate” for [Recovery Policy].

5. Click on [Add a resource to the service], under [Add a new local resource], choose “File System”. Enter a Name, and the appropriate mount point and its device. And check the [Force unmount] checkbox

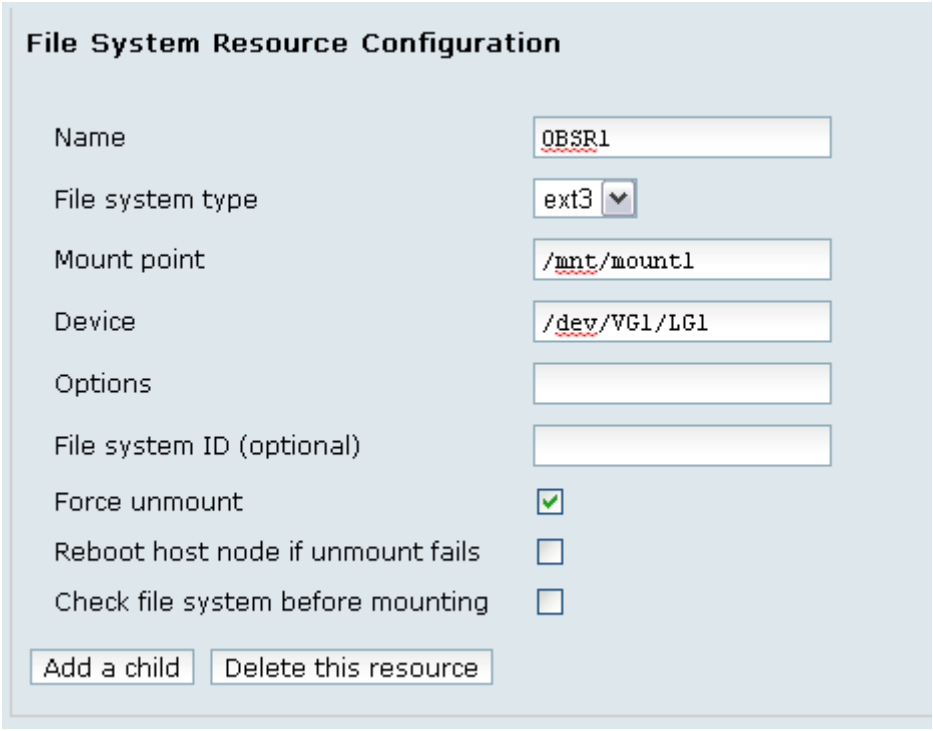

6. Click on [Add a resource to service] again, under [Add a new local resource], choose “IP address”. Enter the desire IP address of the AhsayOBS service, in this example, it is 10.4.0.7

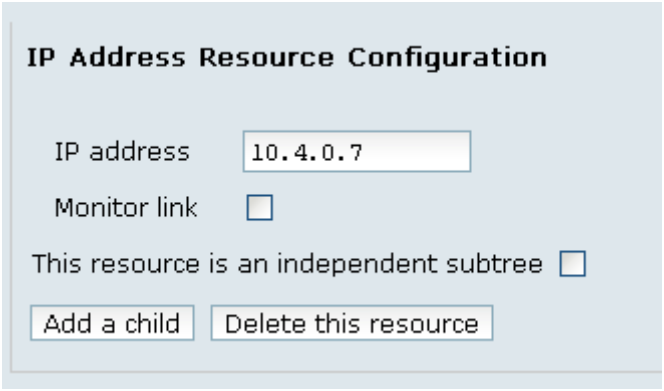

7. Click on [Add a resource to service] again, under [Add a new local resource], choose “Script”. Enter the name of the modified obsr script and its full path.

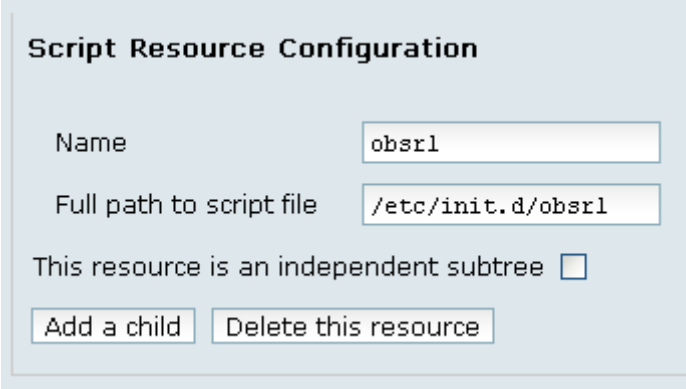

8. Finally, click the [submit] button at the bottom to add the service to the cluster. After the service is added, it should be visible on the services page.

Step2 – Validate your configuration

- Validate if the high availability service is visible in the “Service” page.
- Validate if AhsayOBS service can be started from luci, this can be done with the [Enable this service] option in the “Services” page.
- Validate the configuration by browsing the desire IP address (e.g. http://10.4.0.7:80). This should take you to the AhsayOBS login page.

Important:
You can now activate your AhsayOBS license.

5.5 Setup additional AhsayOBS in the cluster

The procedure in setting up additional AhsayOBS is exactly the same as setting up the first instance, but it must satisfy the following conditions:

- The port numbers must be unique from all other AhsayOBS in the cluster.
- The service name must be unique from all other AhsayOBS in the cluster.
- The server port must be unique from all other AhsayOBS in the cluster.
- Each virtual disk should only be associated with one instance of AhsayOBS, as a virtual drive should only be owned by one Node at any given time.

The steps for setting up additional AhsayOBS in the cluster are summarized as below, where Node1OBSR has already been done throughout the guide.

IP Address	HA Service Name	Installation Location	Mount device	Port Used
10.4.0.7	Node1OBSR	/mnt/mount1/OBSR1	/dev/VG1/LV1	80, 443 & 8014
10.4.0.8	Node2OBSR	/mnt/mount2/OBSR2	/dev/VG2/LV2	70, 8443 & 8015

Please do the followings:

1. Mount device, /dev/VG1/LV2 to /mnt/mount2.
2. Install AhsayOBS at /mnt/mount2/OBSR2.
3. Stop the obsr service and disable obsr service for automatic startup.
4. Rename the script /etc/init.d/obsr with a non-default unique service name. e.g. /etc/init.d/obsr2.
5. Modify the rename script as specify in Appendix A.
6. Deploy the script to the /etc/init.d directory of the other node.
7. Edit server.xml in $OBS_HOME/conf.
8. Edit the startup and shutdown shell under $OBS_HOME/bin.
9. Validate if it can access by go to the Administrator Login Page via web browser. Do not update License at this stage.
10. Setup Node2OBSR as a HA service with IP address being 10.4.0.8.

11. Validate Node2OBSR by invoking http://10.4.0.8:70 and activate its license.

12. Both Node1OBSR and Node2OBSR are now running high availability in the failover cluster.

6 Verify AhsayOBS HA Setup on Linux

After finishing setting up the High Availability Option of AhsayOBS, you must verify the following sections before putting the servers into a production environment.

6.1 Cluster Failover

There are two methods to verify the failover feature of the cluster. However, we would recommend that you should perform both tests to ensure the cluster is working properly.

Method 1 - Manual Failover

A failover event can be simulated manually with the luci interface. Please do the following to trigger a manual failover:

1. Login to luci as administrator

2. Click on the “Cluster” tab and select your cluster.

3. Click on one of the high availability service under the Services Label, this will take you to the Services page.

4. On the Dropdown list that says “Choose a Task”, select “Relocate this service to <Node Name>”, this will simulates a service failover to another Node.

If high availability service, Node1OBSR was running at Node1, you can failover the server to Node2 with luci. During the failover, the file system with OBSR installed will unmount from Node 1 and mount to the same mount point on Node 2, which will execute a start on the defined script.

Method 2 – Auto Failover

To verify the failover ability of the cluster, we should also simulate the event of a server breakdown. Please follow the instructions below, assuming AhsayOBS is running on Node1:

1. Logon to the Node 1.

2. Shutdown the Node 1 with the command:

```
[root] # shutdown now -h
```

3. Login to luci as administrator.

4. Go to the “Services” page.

5. The Status of the service should be: “Running on <Node2>”.

6. Invoke the AhsayOBS Web Admin Console, you are expected to see the standard Administrator page again.

6.2 Other Verification Tests

After you have finished the tests mentioned above, you can verify the AhsayOBS cluster using the AhsayOBM/AhsayACB as you would normally.

- Check that AhsayOBM and AhsayACB can backup and restore.
- Check web restore applet.
- Check whether your branding is correct.
- Check that the email reports for admin and users are correct.
- Check whether data integrity could remain when backup is interrupted by server restarts and temporarily network disconnection.

6.3 Troubleshooting

If there is any issue occurs for the cluster setup, it is advised to check the log located in /var/log/messages or use the clustate command to check the cluster status.

6.4 Additional Information

To start up all cluster related services manually, please perform that in this order:

1. cman
2. clvmd
3. qdiskd
4. rgmanager

To shutdown all cluster related services manually, please perform that in this order:

1. rgmanager
2. qdiskd
3. clvmd
4. cman

7 Server Tuning Recommendations

The performance of AhsayOBS is largely dependent on the computer hardware and operating system installed. Running AhsayOBS with default operating system configuration would produce very good performance throughput. If you would like to fine tune your operating system settings for better performance, please study the suggestions in this chapter.

WARNING:
Please ensure that you understand what you are doing before making any changes to your operating system.

The three major areas of bottleneck affecting the performance and scalability of AhsayOBS are memory, hard disks and network bandwidth.

As with all configuration changes, you must implement the following suggestions one at a time to determine the performance improvements gained. If system performance decreases after making a change, revert back to the original setting.

Memory

Generally, the more memory the server has, the less chance that the operating system needs to swap memory to the disks. Memory swapping to disk is very resource intensive and degrades the performance considerably.

AhsayOBS uses memory for each client connection. The memory required for each connection varies based on the complexity of the backup sets, e.g. the number of files in a single folder. In our experience, each client connection on average consumes approximately 1MB of memory. Therefore, please ensure that there is enough physical memory installed in your servers to support the operating system and the number of simultaneous AhsayOBS client connections that you expect to handle.

On top of this, it is recommended to remove any unused operating system services to free up more memory and CPU resources.

Hard Disks

The performance of AhsayOBS is extremely dependent on disk I/O. Generally, the faster the disk I/0 means the faster AhsayOBS backup/restore operations, and hence, the support of a larger number of simultaneous client connections.

Below are some suggestions to achieve the best I/O performance:

- Use the best hardware available, e.g. 15K rpm SAS Disks.
- Use hardware RAID solutions and avoid using any software-based RAID solutions.
- Consult your hardware provider for advice and best practices for achieving optimum I/O.

Network Bandwidth

Generally, the faster the network connection speed means that backups and restores are faster. In order to maximize the network bandwidth, it is recommended to use Gigabit Ethernet (1000BaseT) NICs and switches. Also, please remember to configure your NICs to operate at the best speed and duplex levels, e.g. hardcode the same speed and duplex settings on the corresponding ports of the network access switch. Do not rely on the [auto-detect] setting.

Besides of this, since majority of operating system are usually tuned for local area network (LAN) access instead of internet access, it is advised to tune the server TCP/IP settings to gain some performance improvement.

8 Upgrade AhsayOBS

8.1 Windows environment

To upgrade the AhsayOBS application, please stop the AhsayOBS service from the Failover Cluster Management Window instead of the Service.msc console.

Stop AhsayOBS Generic Service

1. Go to [Failover Cluster Management]
2. Select your cluster from left hand menu
3. Expand the tree, under the [Services and applications], select the generic service
4. Right click on it and select "Take this service or application offline" to stop the service.

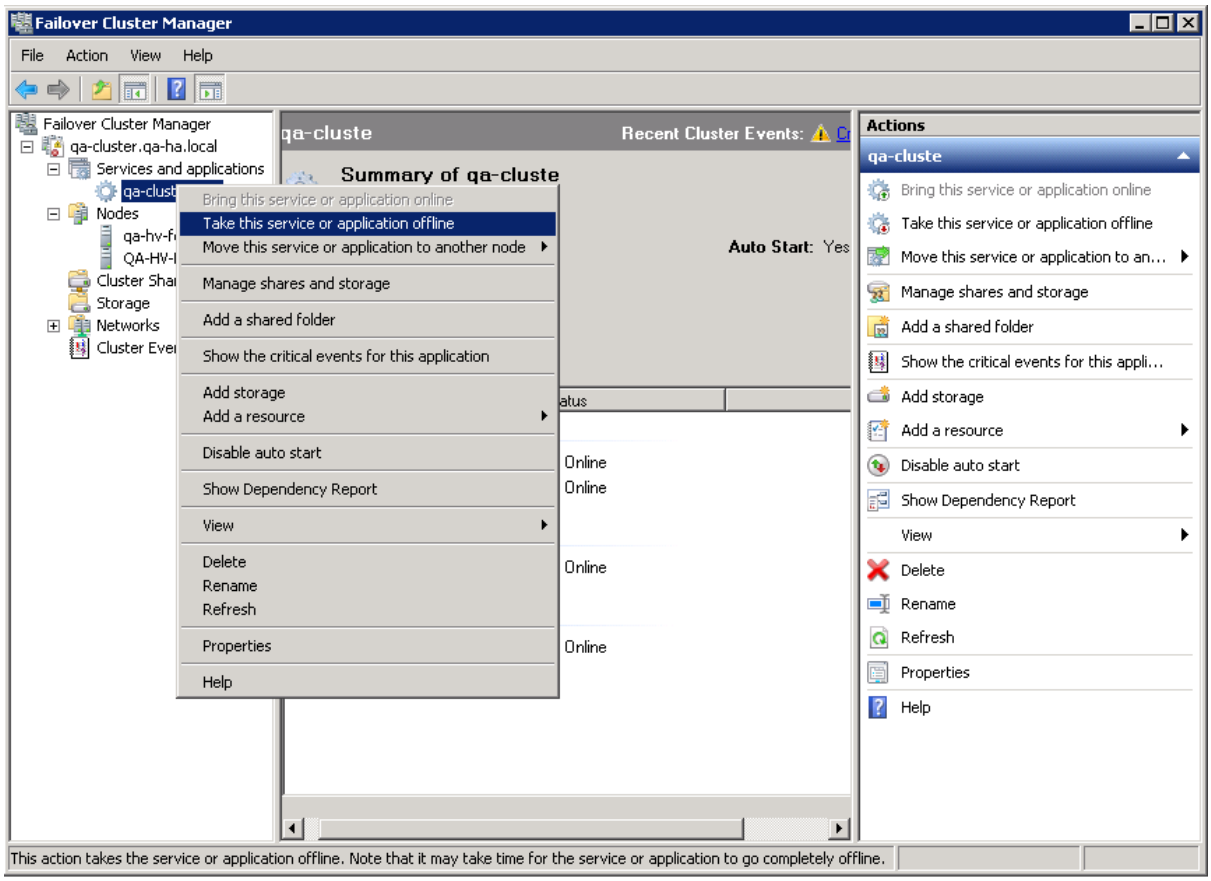

However, by doing this the Shared Cluster Disk of the AhsayOBS service will also go offline. Therefore, we need to bring back the storage online only by following steps:

1. Go to [Failover Cluster Management]

2. Select your cluster from left hand menu
3. Expand the tree, under the [Services and applications], select the generic service
4. Select the generic service.
5. In the middle menu, right click the disk, and click "Bring this resource online"
6. After the disk is online, it is available to one of the cluster machine.

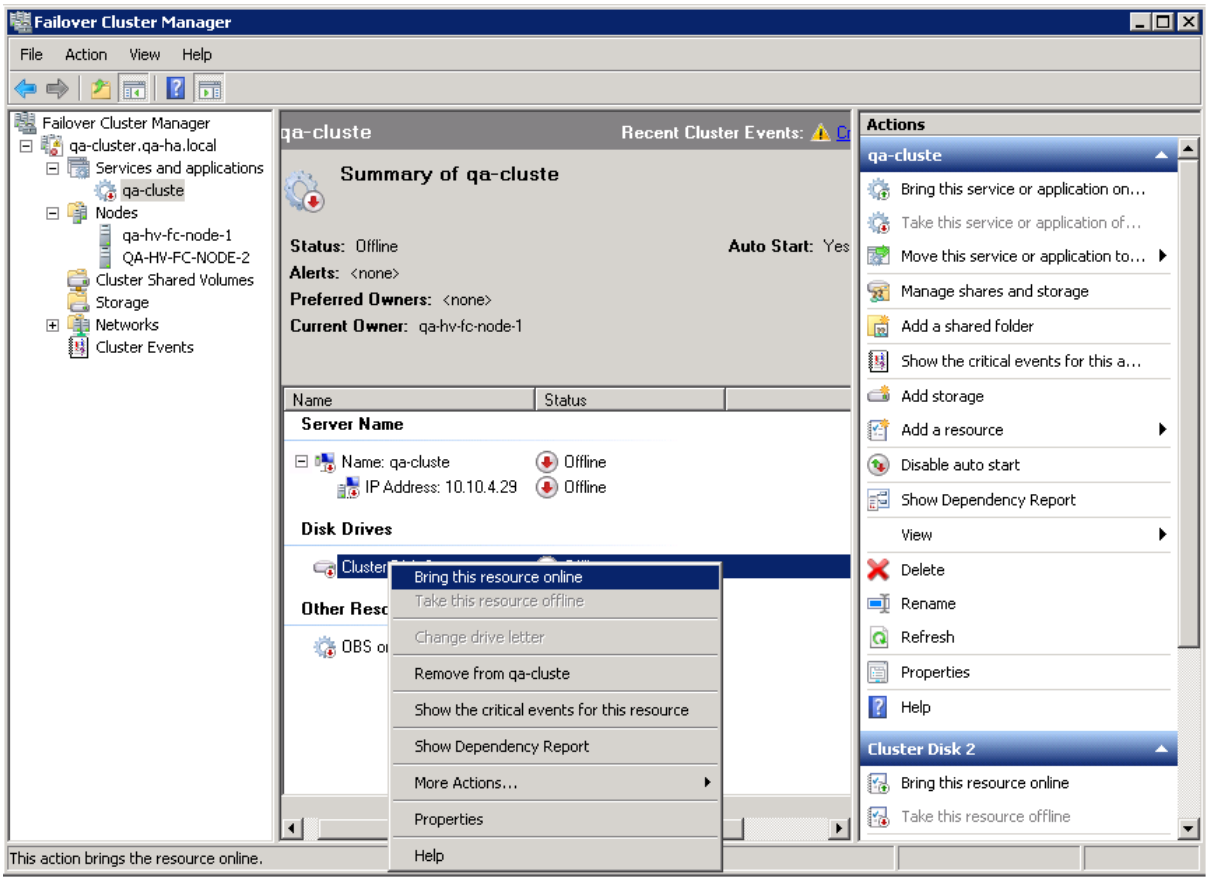

After the disk is online, perform the upgrade on AhsayOBS.

Start AhsayOBS Generic Service

1. Go to [Failover Cluster Management]
2. Select your cluster from left hand menu
3. Expand the tree, under the [Services and applications], select the generic service.
4. Right click on it and select "Bring this service or application online" to start up the service.

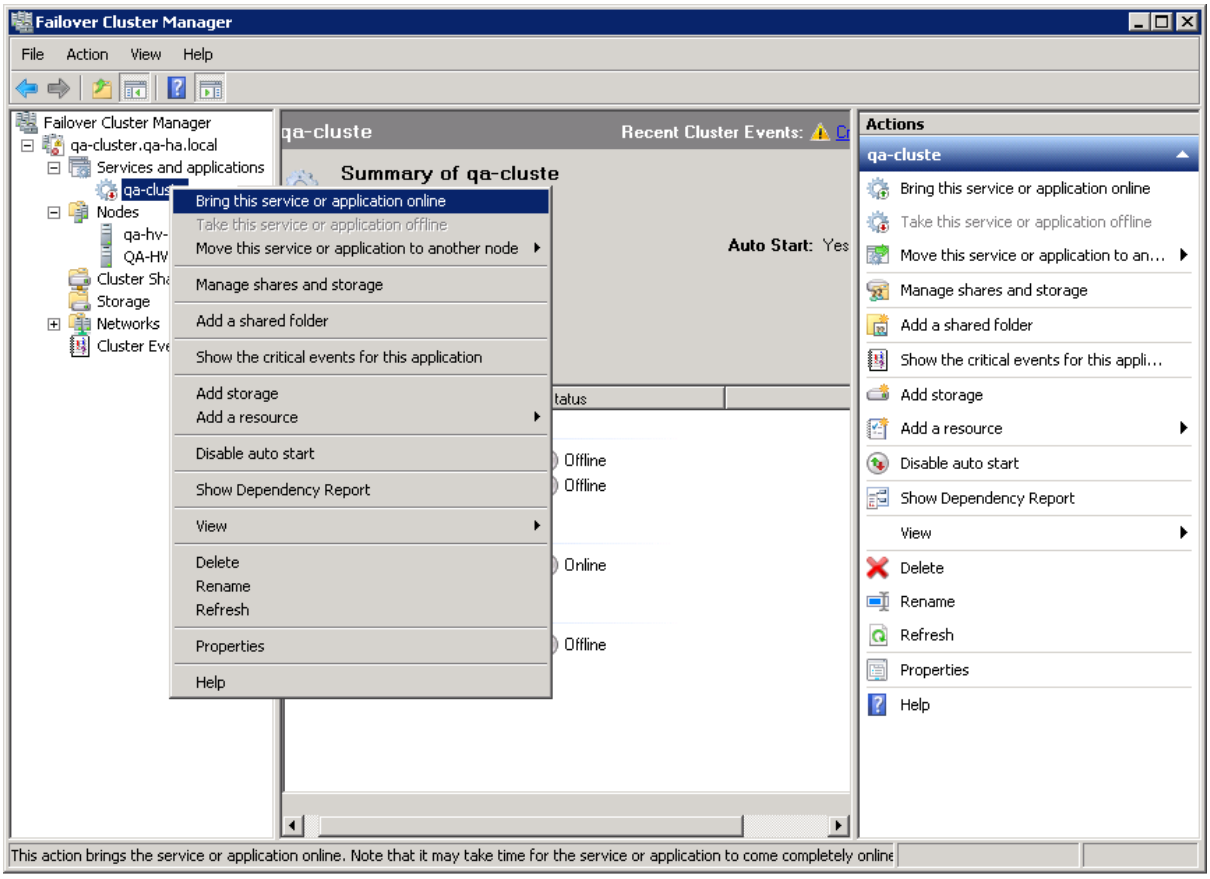

Note:
Sometimes, there is a java.exe process leaves behind on the machine after shutting down the generic service. User may require to kill the java process manually.

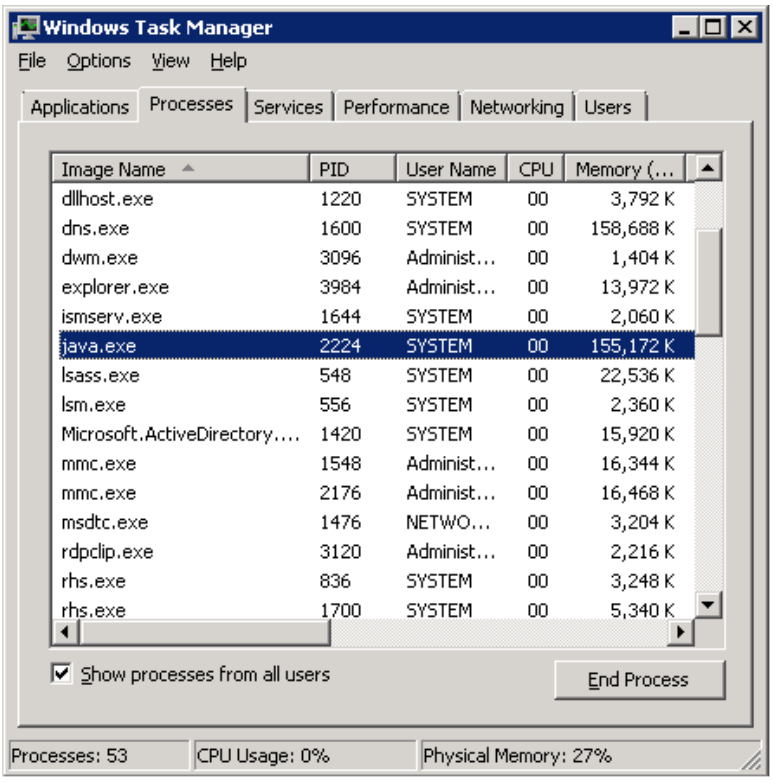

9 References

- Dell PowerVault MD3000 with Windows Server® Failover Clusters Hardware Guide, available at:

  http://docs.us.dell.com/support/edocs/systems/clusters/se600w/en/Hardware_IT/pdf/it_en.zip

- Dell PowerVault MD3000 with Windows Server® Failover Clusters Software Guide, available at:

  http://docs.us.dell.com/support/edocs/systems/clusters/se600w/en/Software_it/Win08/pdf/swit_doc.zip

- Microsoft Quorum configuration in a failover cluster for Windows 2008, available at:

  http://technet.microsoft.com/en-us/library/cc770620(WS.10).aspx

Appendix

Appendix A Product Documentations

Please visit this link for the documentations of Ahsay Products.

Appendix B Sample Script for AhsayOBS HA on Linux

Modify obsr script to run high availability, suppose the unique service name is **obsr1**.

```
#!/bin/bash
#
# obs: Startup AhsayOBS and AhsayRPS (Ahsay Offsite Backup Server and
Replication Server)
#
# chkconfig: 2345 95 95
#
# description: Running AhsayOBS and AhsayRPS on system startup
#
#

# [Ahsay Offsite Backup Server and Replication Server] (version 6.0.3.1)

# Source function library
. /etc/init.d/functions

OBSR_HOME=/mnt/mount1/OBSR1 ← This is your OBS_HOME
OBSR_USER=root
OBSR_PIDFILE=/var/run/obsr1.pid ← This will be your service name

export OBSR_HOME OBSR_USER OBSR_PIDFILE

RETVAL=0

case "$1" in
 start)
    echo ""
    echo "Starting up [ Ahsay Offsite Backup Server and Replication Server ]"
    if [ root = ${OBSR_USER} ];
    then
      /bin/sh ${OBSR_HOME}/bin/startup.sh
    else
      su ${OBSR_USER} -c "/bin/sh ${OBSR_HOME}/bin/startup.sh"
    fi

    sleep 5
    echo "[ Ahsay Offsite Backup Server and Replication Server ] is running"
    ;;

 stop)
    echo ""
    echo "Shutting down [ Ahsay Offsite Backup Server and Replication Server ]"
    if [ root = ${OBSR_USER} ] ;
    then
      /bin/sh ${OBSR_HOME}/bin/shutdown.sh
```

```
    else
      su ${OBSR_USER} -c "/bin/sh ${OBSR_HOME}/bin/shutdown.sh"
    fi

    sleep 5
    echo "[ Ahsay Offsite Backup Server and Replication Server ] is stopped"
    ;;
 status)
    echo ""
    OBSR_NAME=$0
    if [ -f ${OBSR_PIDFILE} ]; then
      status -p ${OBSR_PIDFILE} ${OBSR_NAME}
    else
      status ${OBSR_NAME}
    fi
    RETVAL=$?
  ;;

    *)
    echo $"usage: $0 {start|stop|status}";;

esac

exit $RETVAL ← Replaced exit 1
```