Technical white paper

Adding a cluster node to an HP StoreEasy 3000

Table of contents

Click here to verify the latest version of this document

Introduction

Today's 24x7 economy demands the availability and reliability of data to be competitive in the market. Small and medium businesses (SMBs) cannot afford to absorb the high costs and loss of productivity due to loss of critical business data. For example, losing even a single order or single email communication from a customer can seriously harm business. Also, with the growing business needs there is a significant amount of data explosion which calls for intensive CPU and IO processing demands. This stresses the need to create a fault tolerant and highly available environment.

The HP StoreEasy 3000 Storage products are highly available products using failover clustering consisting of two or more nodes. By presenting LUNs from supported storage arrays to HP StoreEasy 3000, you can offer highly available SMB and NFS file shares to clients.

Adding nodes to a failover cluster increases system redundancy, improves load balancing and enables highly available services such as file shares, iSCSI targets, and more. Clients see only a brief interruption of service as the cluster resource transitions from one server node to the other when a failure occurs. Windows® Storage Server 2012 offers improved failover cluster support for applications such as Hyper-V, Microsoft® SQL servers, Microsoft Exchange servers and file servers. The cluster scalability has been extended up to 64 nodes in a cluster.

This document provides details regarding the configuration of additional node to the cluster and recommends best practices.

Cluster terminologies

Failover Cluster is a group of independent computers or nodes that work together to increase the availability of applications and services.

Nodes are clustered servers connected by physical cables and software.

Cluster Resources are hardware and software components that are managed by the cluster service. It includes IP address, cluster name, cluster quorum disk, physical disk, file servers, applications etc.

Heartbeat network is a private network to track the state of each cluster node. Each node sends out periodic messages to the other nodes; these messages are called heartbeats. If a node stops sending heartbeats, the cluster service fails over any resources that the node owns to another node.

Quorum disk is the shared storage used by the cluster nodes to coordinate the internal cluster state. The Quorum disk maintains data integrity by storing the most current version of the cluster database.

Cluster Shared Volumes (CSVs) in a Windows Server® 2012 failover cluster allow multiple nodes in the cluster to simultaneously have read-write access to the same LUN (disk) that is provisioned as an NTFS volume.

Adding node to cluster

Pre-requisites

This paper assumes that you already have HP StoreEasy 3000 Storage systems installed as a two-node failover cluster.

Complete the initial installation of the new storage server by following the HP StoreEasy Quick Start Guide. At the time of installation there would be two options to either deploy it as standalone system or Windows Failover Cluster. Choose to deploy it as standalone system at this stage. We will add it to the Windows failover cluster later.

Configure Networking

Configure a switch or use an existing switch with a dedicated heartbeat VLAN which would be used for the cluster heartbeat network. Disconnect the crossover cable between the first node and second node and connect it to the switch and assign IP addresses. Connect Port 4 of the new node to the next available port on the same switch and assign the port to the heartbeat VLAN. The nodes would use this network as heartbeat network to check if each of the nodes is alive. The heartbeat IP address needs to be in the same subnet as the other cluster node heartbeat network connections. It is recommended to isolate the heartbeat network to a separate VLAN or subnet in order to eliminate unnecessary traffic from the heartbeat network adapter which is set for internal cluster communications only.

Configure the remaining network adapters for use in your network infrastructure and one of the adapters needs to be configured which would provide a route to the domain controller since we will need to join the server to an Active Directory domain. If you are adding more than one adapter per server to your network infrastructure, each adapter should be on a different subnet.

Provision shared storage—Only Quorum disk

A Windows failover cluster requires shared storage from the storage array. All storage provisioning for HP StoreEasy 3000 systems are done on the particular array used for storage. Consult the documentation for your particular array to perform the necessary tasks involved in presenting LUNs to the HP StoreEasy 3000. In general, you should follow the array's guidelines for providing storage to Windows Storage Server 2012. This will also likely involve tasks such as cabling, configuring ports, zoning, and configuring MPIO.

When adding another node to an existing cluster there would be LUNs that have been already created and presented to the existing nodes. Once the connection to the array has been set, present only the existing Quorum disk to the new node. The other existing volumes should not be made accessible to the new node after it has joined the cluster. This is to avoid data corruption which may be caused presenting LUNs to non-clustered system. It can be viewed from Server Manager -> File and Storage Services -> Volumes or use the Windows Disk Management utility.

After the node has joined the cluster, the LUNs may be presented to the new node. Move the physical disk resources over to the new node to confirm functionality.

Run Microsoft Windows Update

HP highly recommends that you run Microsoft Windows update to identify, review, and install the latest, applicable, critical security updates on the storage system. If there is a Service Release released by HP which is already running on the existing nodes in the cluster then it has to be applied on the new node as well.

Join the cluster to domain

The new node needs to be added to the same domain to which the existing cluster nodes are previously joined and this can be done either using ICT or Server Manager.

Add Node

From one of the existing nodes in the cluster, select the Server Manager icon in the desktop taskbar. Then select "**Tools**" -> "**Failover Cluster Manager**".

Existing cluster would be displayed in the left navigation pane. Right-click the cluster name and select "**Add Node**" or select the cluster name to view the Add Node option under "**Configure**" in the main viewing pane. It would open up the "**Add Node Wizard**".

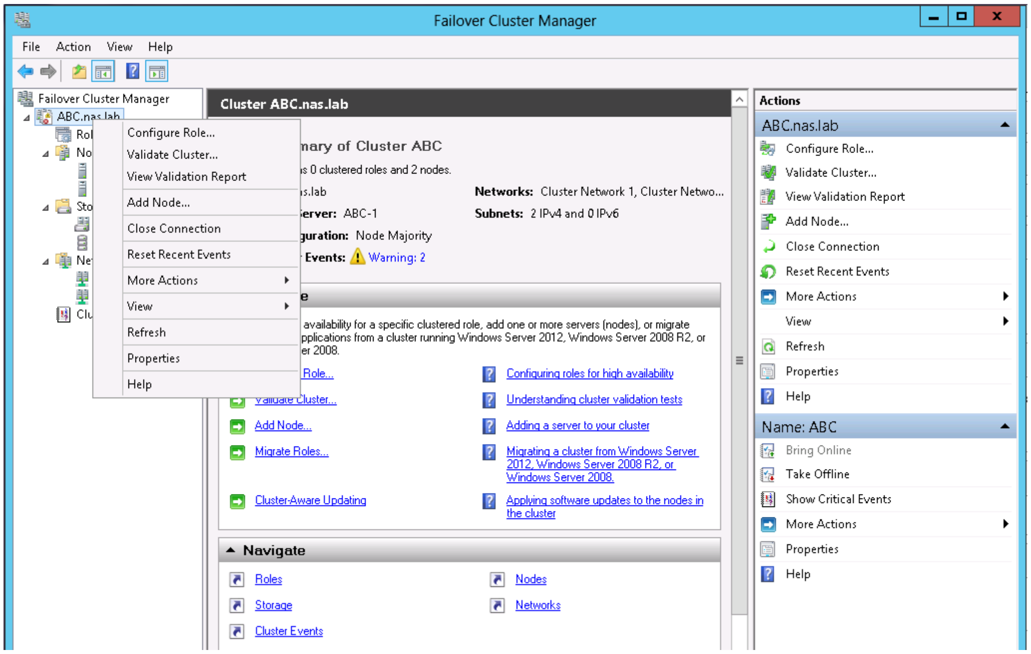

Enter the name of the node to be added to the cluster and click “**Add**”.

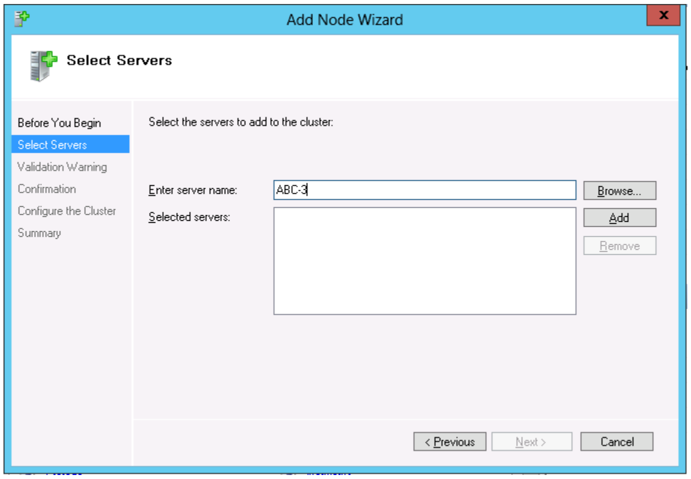

Select the servers to add to the cluster:

Enter server name: ABC-3

Selected servers:

In the confirmation page check the option “**Add all eligible storage to the cluster**”.

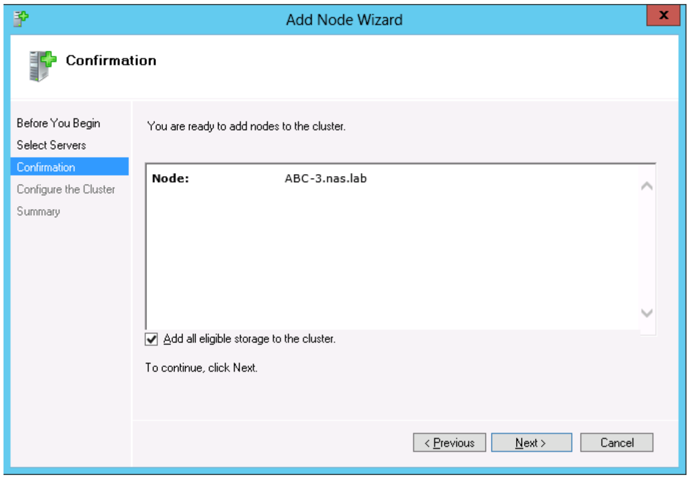

You are ready to add nodes to the cluster.

Node: ABC-3.nas.lab

☑ Add all eligible storage to the cluster.

To continue, click Next.

Complete the wizard and view the report to ensure there are no failures and the operation is successful.

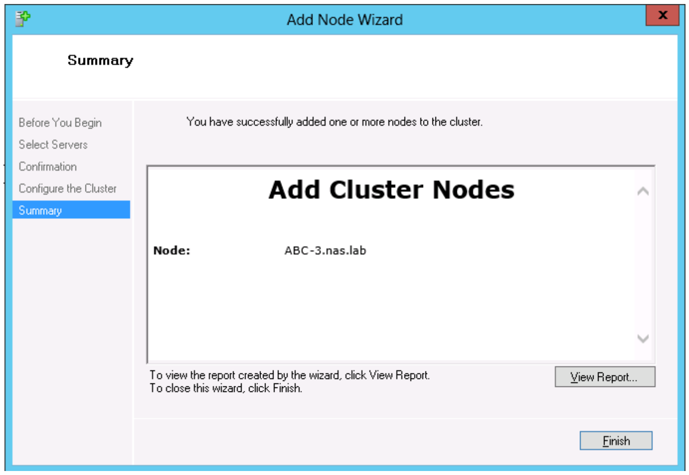

You have successfully added one or more nodes to the cluster.

Add Cluster Nodes

Node: ABC-3.nas.lab

To view the report created by the wizard, click View Report.
To close this wizard, click Finish.

Troubleshooting tips

- No shared disk for the Quorum disk is found. A shared disk with a NTFS partition at least 512 MB in size should be presented to all the nodes in the cluster.
- Shared storage presented from the storage array but not formatted with a NTFS partition can throw errors during cluster creation. Use the Windows Disk Management utility to configure shared disk resources. Verify that all shared disks are formatted as NTFS and are designated as Basic.
- Issues connecting to the nodes using the hostname. Verify name resolution, ping each node from a client using the node's machine name instead of its IP address and DNS server is functioning properly.
- The user has no permission to do certain cluster operations. This user account will need to be granted administrator privileges.
- Use of DHCP addresses for network connections can lead to network validation errors at times. All network adapters must be configured with static IP addresses in a cluster configuration.
- Errors appear on a network adapter that is not configured or does not have an active link. If the network adapter is not going to be used it may be disabled.

Best Practice recommendations

Hardware redundancy

It is important to eliminate single point of failure by maintaining redundancy as much as possible in terms of hardware and infrastructure.

- It is recommended that cluster configurations be deployed with dual data paths for high availability. Consider connecting at least two single Host Bus Adapters (HBAs) to the storage array which would provide an alternate path in case of failures.
- Introduce redundancy at the network switch level by having two switches so that the cluster nodes can withstand a switch failure and communication is not disrupted.
- It is recommended to have multiple network cards per StoreEasy 3000 node.

Software configuration

- All the nodes in the cluster should be running the same version of service packs and updated with security patches and hotfixes using Windows Update. In the event of HP releasing Service Release for StoreEasy 3000, it needs to be installed on all the nodes.
- Ensure DNS is configured in the network and the cluster nodes are able to resolve names using DNS.
- In case of installing or removing any roles or features it needs to be replicated on each node in the cluster.
- Plan a maintenance window to run system health check regularly. At Run type "**perfmon/report**" to invoke system health check. It would do basic system checks and report the results of the tests to the user. In case of any issues you need to take corrective actions.
- The Best Practices Analyzer (BPA) is a server management tool which would scan and provide guidelines to configure a server as defined by experts. The BPA GUI can be used to start BPA scan which is available under Tasks menu of the Best Practices Analyzer for each of the Server Manager role. Click "**Start BPA Scan**".

Network settings and configuration

A cluster uses at least two network connections on each node:

- The private cluster interconnect or "heartbeat" crossover cable connects to one of the network ports on each cluster node. In more than two node deployments, a private VLAN on a switch or hub is required for the cluster interconnect.
- The public client network subnet connects to the remaining network ports on each cluster node.

The heartbeat network which checks the health status of the cluster needs to be isolated from the regular traffic and as mentioned earlier during the configuration step it is recommended to isolate the heartbeat network to a separate VLAN or subnet. Depending on the latency of your network you could increase the interval and increase/decrease the threshold.

The following procedures are best practices provided by Microsoft and should be configured on the private network adapter.

- On the General tab of the private network adapter, ensure that only TCP/IP is selected.
- Ensure that the Register this connection's address in DNS is not selected in the DNS tab under advanced settings for Internet Protocol (TCP/IP) Properties.
- In all cases, set static IP addresses for the private network connector.

It is strongly recommended to set static IP addresses for all network adapters in the cluster, both private and public. If IP addresses are obtained through DHCP, access to cluster nodes could become unavailable if the DHCP server goes down. If DHCP must be used for the public network adapter, use long lease periods to assure that the dynamically assigned lease address remains valid even if the DHCP service is temporarily lost.

It is important there are no two or more NICs on the same node that are configured to be on the same subnet. This is because the cluster network driver uses the subnet to identify networks and will use the first one detected and ignore any other NICs configured on the same subnet on the same node. The cluster validation process will register a Warning if any network interfaces in a cluster node are configured to be on the same network.

HP recommends changing the names of the network connections for clarity and easy identification as per its role. For example, "Cluster interconnect" for the private network and "Public connection" for the public network.

It is recommended to dedicate a NIC for Hyper-V network which would be required for live migrations. If you use a network for iSCSI (storage), do not use it for network communication in the cluster. For more information on Hyper-V networking check technet.microsoft.com/en-us/library/ff428137(WS.10).aspx and for iSCSI refer technet.microsoft.com/en-us/library/jj612869.aspx.

Quorum configuration

The cluster quorum configuration impacts the high availability of the cluster and the roles hosted on that cluster. It ensures the cluster functions in case of failover or membership change. In the event of split cluster where one subset of nodes cannot communicate with another subset of nodes, the quorum configuration ensures only one set of subset runs as a cluster to avoid data corruptions. Dedicate a separate disk resource for a Quorum disk because the failure of the Quorum disk would cause the entire cluster to fail. HP strongly recommends that the disk resource be a RAID 1 configuration.

The user would not be asked to choose the quorum model during cluster creation instead it is automatically selected based on the number of nodes and cluster resources. However, changes can be made to the quorum model if required. There are basically 4 quorum types:

- Node Majority—Recommended in case of odd number of nodes. Can sustain failures of half the nodes minus one.
- Node and Disk Majority—Recommended when there are even number of nodes. Can sustain failures of half the nodes if the disk witness remains online.
- Node and File Share Majority—Same as Node and Disk Majority except that here the File Share is the witness resource.
- No Majority or Disk only—Not recommended since in this case the disk is the single point of failure. It can survive failure of all nodes except one if the disk is online.

Right click on the cluster listed in Failover Cluster Manager, click “**More Actions**” and select “**Configure Cluster Quorum Settings**” to change the quorum configuration. For more information on configuring and managing the quorum, see technet.microsoft.com/en-us/library/jj612870.aspx.

Failover/Failback Policies

The movement of applications in a cluster in the event of planned or unplanned failure of the node or resources is referred to as failover. In clusters containing more than two nodes, additional fail over rules can be applied. For instance, groups can be configured to failover different nodes to balance the additional work load imposed by the failed node. Nodes can be excluded from the possible owners list to prevent a resource from coming online on a particular node. Lastly, the preferred owners list can be ordered, to provide an ordered list of failover nodes. The failover of resources can be controlled with in a multinode cluster to provide a controlled balanced failover methodology that balances the increased workload.

Because operating environments differ, the administrator must indicate whether the system will automatically fail the resources (organized by resource groups) back to their original node or will leave the resources failed over, waiting for the resources to be moved back manually. The failback policy must be set before the failover occurs so that failback works as intended. The failback policy is not set by default and unless you manually configure the node to failback after failover the applications would continue to run on the alternate node after the failed node comes back online.

To configure failover and failback settings for the clustered service or application, open the Failover Cluster Manager and connect to the cluster. Click on “**Roles**” and right-click on the role and select “**Properties**”. Click on the Failover tab and configure the failover and failback policy.

The failover and failback policies need to be tested prior to moving to production to ensure the nodes can take over the services without any issues.

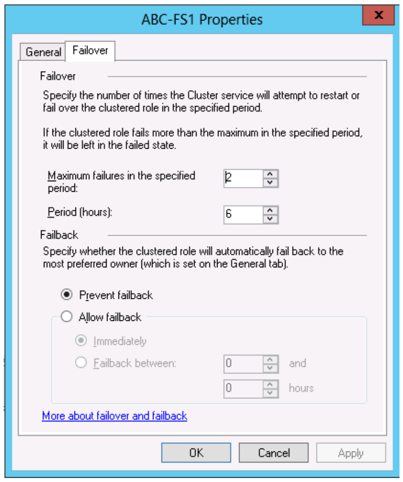
ABC-FS1 Properties

General | Failover

Failover

Specify the number of times the Cluster service will attempt to restart or fail over the clustered role in the specified period.

If the clustered role fails more than the maximum in the specified period, it will be left in the failed state.

Maximum failures in the specified period: 2

Period (hours): 6

Failback

Specify whether the clustered role will automatically fail back to the most preferred owner (which is set on the General tab).

● Prevent failback

○ Allow failback

● Immediately

○ Failback between: 0 and 0 hours

More about failover and failback

OK | Cancel | Apply

Cluster-Aware Updating

Cluster-Aware Updating (CAU) is a reliable, automated, and integrated update feature that allows system administrators to update clustered nodes with little or no loss in availability during the update process. It would first scan the nodes to find which patches are needed and then identify the node with lowest workload. It would then move the workloads, put the node into maintenance mode and calls the Windows Update Agent to patch the server and verify the patch installation is successful. It then moves on the other nodes and repeats the same steps on all the remaining nodes.

CAU can coordinate the complete cluster updating operation in two modes—Self-updating mode and Remote-updating mode. In self-updating mode, CAU can update the failover cluster by using a fully automated, end-to-end updating process. For remote-updating mode a remote computer that is running Windows Storage Server 2012 R2, which is called an Update Coordinator, is configured with the CAU tools and the administrator triggers an on-demand Updating Run by using a default or custom Updating Run profile.

To use Cluster-Aware Updating, from Failover Cluster Manager, right click on the cluster and from "**More Actions**" select "**Cluster-Aware Updating**". The CAU clustered role is not enabled by default so you will first need to click the "**Configure self-updating options**" action on the main console under Cluster Actions. This role is required to start a Self-Updating Run from a cluster node. You would be able to set the frequency for self-updating. The option "**Preview updated for this cluster**" would generate list of updates which can be reviewed. You could click on "**Apply updates to this cluster**" if you want the updates to be applied at any point of time.

For more information on Cluster-Aware Updating, see technet.microsoft.com/library/hh831694.aspx.

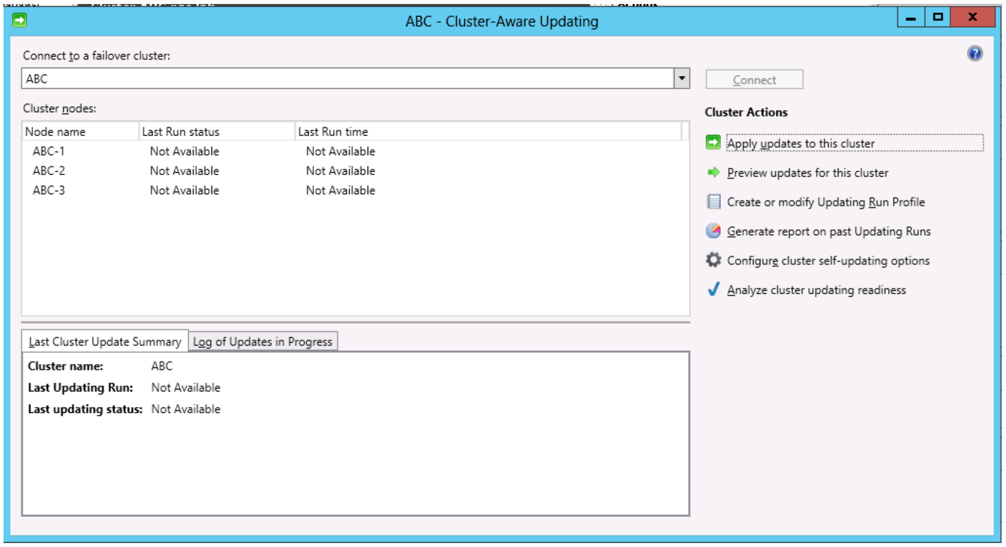
ABC - Cluster-Aware Updating

Connect to a failover cluster: ABC

Connect

Cluster nodes:

Node name	Last Run status	Last Run time
ABC-1	Not Available	Not Available
ABC-2	Not Available	Not Available
ABC-3	Not Available	Not Available

Cluster Actions

- Apply updates to this cluster
- Preview updates for this cluster
- Create or modify Updating Run Profile
- Generate report on past Updating Runs
- Configure cluster self-updating options
- Analyze cluster updating readiness

Last Cluster Update Summary | Log of Updates in Progress

Cluster name: ABC

Last Updating Run: Not Available

Last updating status: Not Available

References

HP StoreEasy 3000 Storage
hp.com/go/StoreEasy

Failover Clustering Overview
technet.microsoft.com/en-us/library/hh831579.aspx

What's New in Failover Clustering in Windows Server 2012?
technet.microsoft.com/en-us/library/hh831414.aspx

Configure Failover and Failback settings
technet.microsoft.com/en-us/library/dd197473(v=ws.10).aspx
support.microsoft.com/kb/197047

Learn more at
hp.com/go/StoreEasy

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document