HITACHI
Inspire the Next

SMB 3 Transparent Failover for Hitachi NAS Platform 4000 Series

TECHNICAL PAPER

Tech Note

By Leo Ngyuen

November 2016

Hitachi Data Systems

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Contents

SMB 3 Transparent Failover for Hitachi NAS Platform 4000 Series

Tech Note

This tech note documents the results of transparent failover testing performed in the Solution Engineering labs at Hitachi Data Systems to verify whether Hitachi NAS Platform 4000 series (HNAS 4000) supports transparent failover for Microsoft® Hyper-V® for virtual machine workloads hosted on a NAS Platform file share.

A file share hosted by the Hitachi NAS Platform cluster successfully supported transparent failover in a Microsoft Hyper-V environment.

The testing used two server blades on Hitachi Compute Blade 500 to host the Hyper-V environment. These server blades were configured as a failover cluster. Two Hitachi NAS Platform 4080 systems were configured as a cluster and managed by an SMU 400. The NAS Platform cluster was connected to Hitachi Virtual Storage Platform G600 (VSP G600).

Server Message Block (SMB) transparent failover is one of the key features introduced in SMB 3.0. With SMB Transparent Failover, you can configure file shares in Microsoft Windows Server® using failover clustering to be continuously available. Using continuously available file shares enables you to perform hardware or software maintenance on any cluster node on Hitachi NAS Platform 4080 without interrupting the server applications that are storing their data files on these file shares.

This document provides the following:

- A proof point of the basic functionality of this solution
- High level technical reference to consider this solution in your environment

This document **does not** cover the following:

- Sizing information
- Best practice
- Implementation details

For implementation details, contact your Hitachi Data Systems representative.

Note — Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Solution Components

These are the hardware and software components used to validate the test environment.

Hardware Components

Table 1 lists the information about the hardware components used in the test environment.

Table 1. Hardware Components

Hardware	Description	Version	Quantity
Hitachi NAS Platform 4080	Network attached storage (NAS) platform: ■ 4 × 10 GbE Ethernet ports ■ 4 × 8 Gb/sec Fibre Channel ports ■ 2 × 10 GbE cluster interconnect ports	12.5.4037.00	2
System Management Unit software 400	■ 2 × 10/100/1000 Mb/sec Ethernet ports	12.5.4037.00	1
Hitachi Virtual Storage Platform G600	Enterprise storage technology with flash-optimized performance: ■ Dual controllers ■ 16 × 8 Gb/sec Fibre Channel ports ■ 256 GB cache memory ■ 128 × 1.2 TB, 10k RPM, SAS disks	83-01-01-40	1
Hitachi Compute Blade 500 (CB 500)	Sever blade chassis: ■ Up to 8 server blades ■ 2 management modules ■ 6 cooling fan modules ■ 4 power supply modules ■ 2 Brocade 5460 Fibre Channel switch modules ■ 2 Brocade VDX 6746 DCB 10 GbE switch modules	A0160-E-8004	1

Table 1. Hardware Components (Continued)

Hardware	Description	Version	Quantity
520H B2 server blade	Half-size server blade: ■ 2 × 12-core Intel Xeon E5-2697 v2 processors at 2.70 GHz ■ 256 GB RAM ■ 16 × 16 DIMMs ■ 2 hot-swappable 2.5 inch SAS drives ■ 1 Emulex 10 GbE onboard CNA ■ 1 Hitachi 16 Gb/sec Fibre Channel mezzanine card	04-05/10-10	2
Brocade 6740	24-port 10 GbE switch	2.0.1b	2
Brocade 6510	48-port 8-16 Gb/sec Fibre Channel switch	7.0.1.a	2

Software Components

Table 2 lists the software used in the test environment.

Table 2. Software Components

Software	Version
Hitachi Storage Virtualization Operating System with Hitachi Dynamic Provisioning	VSP G600 microcode dependent
Hitachi Storage Navigator	VSP G600 microcode dependent
Microsoft Windows Server	2012 R2 Datacenter
	2012 R2 Standard
Microsoft System Center	2012 R2
Microsoft System Center Virtual Machine Manager	2012 R2
Microsoft SQL Server®	2012

Note the following:

- Hitachi Dynamic Provisioning was used on Hitachi Virtual Storage Platform G600.
- Microsoft Windows Server 2012 R2 Datacenter was installed on the Microsoft Hyper-V hosts.
- Microsoft Windows Server 2012 R2 Standard was installed on all virtual machines.
- Microsoft System Center 2012 R2 was installed in the test environment, but not used for the initial validation testing.
- Microsoft SQL Server 2012 was installed in the environment to host the required System Center databases.

The following roles and features were installed on both of the server blades:

- Hyper-V role
- Multipath I/O
- Failover clustering

Test Environment Configuration

The test environment consisted of the following:

- Two server blades configured as a failover cluster with the Microsoft Hyper-V role enabled.
- Hitachi Virtual Storage Platform G600 with one dynamic provisioning pool configured to hold configuration stores and VHD files for the infrastructure virtual machines and web server virtual machines.
- Two Hitachi NAS Platform 4080 systems configured as a cluster, managed with an SMU 400.
- Hitachi Virtual Storage Platform G600 for functionality and performance testing with the following:
  - 1 dynamic provision pool
  - 8 parity groups configured as RAID-6 (6D+2P)
  - 16 LUNs presented to the Hitachi NAS Platform cluster
- Each of the 16 LUNs from the Virtual Storage Platform G600 storage array was configured as a 1.04 TB system drive on Hitachi NAS Platform. One storage pool was created with a single stripeset, consisting of 16 system drives. A 20 TB file system was configured as a storage pool. The file system was assigned to an EVS configured on the Hitachi NAS Platform 4080 cluster. A CIFS server was configured and a CIF share was created for the file system.

Figure 1 shows the high level architecture of the test environment. Infrastructure Fibre Channel switches and Ethernet switches are not shown. All Fibre Channel connections were made through a pair of Brocade 6510 switches.

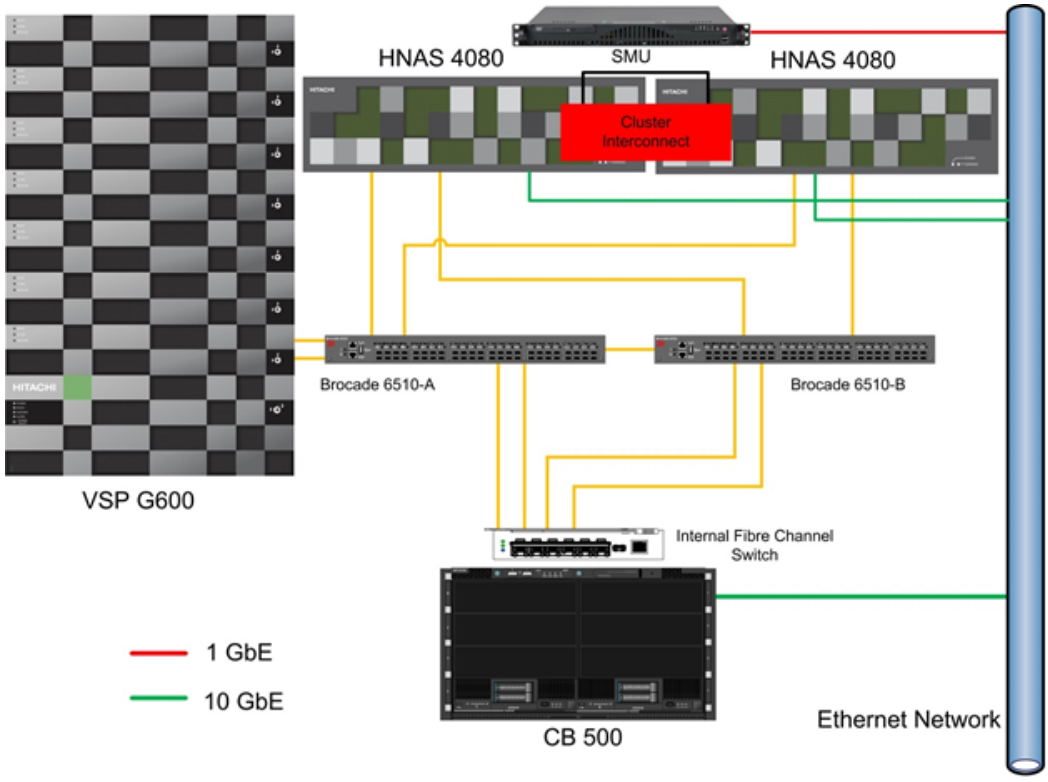

Figure 1

Storage Configuration

One Hitachi Virtual Storage Platform G600 storage system was used for the test environment:

- Attached to Hitachi Compute Blade 500
- Hosted VHDX files and configuration files for infrastructure virtual machines and Web Server virtual machines.

Hitachi Virtual Storage Platform G600 Configuration

Two dynamic provisioning pools were configured on Hitachi Virtual Storage Platform G600. The configuration details for the pool are in Table 3.

Table 3. Virtual Storage Platform G600 HDP Pool Configuration

Dynamic Provisioning Pool	Dynamic Provisioning Pool RAID Configuration	Number of Drives	Drive Capacity	Pool Capacity	Use
HYPER-V-1	RAID-6 (6D+2P)	64	1.2 TB, 10k RPM SAS	23.99 TB	Functionality performance testing

Three LDEVs were allocated from the dynamic provisioning pools and mapped to storage ports as shown in Table 4.

Table 4. Virtual Storage Platform G600 LDEV Configuration

Dynamic Provisioning Pool	LDEV	LUN	Size (GB)	Use	Storage Port
HYPER-V-1	00:00:82	0	1	Quorum	5A, 6A
	00:00:83	1	500	Cluster Shared Volume	5A, 6A

The LUNs were presented to two Hitachi Compute Blade 500 server blades that were configured in a failover cluster and used for the following:

- **LUN 0** — Configured as the quorum disk for the failover cluster
- **LUN 1** — Configured as a cluster shared volume and used for the configuration store and VHDs for the environment infrastructure virtual machines

Hitachi Virtual Storage Platform G600 Storage Presented to Hitachi NAS Platform 4080 Configuration

Storage allocated to the Hitachi NAS Platform 4080 cluster consisted of eight RAID groups. Each RAID group consisted of four 1.2 TB 10k RPM SAS drives, configured as RAID-6 (6D+2P).

One LDEV was created in each RAID group. Each LDEV was presented to the Hitachi NAS Platform 4080 cluster as a LUN. Each LUN was configured as a system drive.

One storage pool was created on the Hitachi NAS Platform 4080 cluster consisting of 16 system drives. Table 5 shows the configuration.

Table 5. Hitachi Virtual Storage Platform G600 Storage Presented to Hitachi NAS Platform 4080 Configuration

LDEV	LUN	Storage Ports	Hitachi NAS Platform Storage Pool
00:00:6A	0	5A,6A	HNAS-FS-1
00:00:6B	1	5A,6A	HNAS-FS-1
00:00:6C	2	5A,6A	HNAS-FS-1
00:00:6D	2	5A,6A	HNAS-FS-1
00:00:6E	4	5A,6A	HNAS-FS-1
00:00:6F	5	5A,6A	HNAS-FS-1
00:00:70	6	5A,6A	HNAS-FS-1
00:00:71	7	5A,6A	HNAS-FS-1
00:00:72	8	5A,6A	HNAS-FS-1
00:00:73	9	5A,6A	HNAS-FS-1
00:00:74	10	5A,6A	HNAS-FS-1
00:00:75	11	5A,6A	HNAS-FS-1
00:00:74	12	5A,6A	HNAS-FS-1
00:00:73	13	5A,6A	HNAS-FS-1
00:00:74	14	5A,6A	HNAS-FS-1
00:00:75	15	5A,6A	HNAS-FS-1

Hitachi NAS Platform 4080 Configuration

The Hitachi NAS Platform configuration is a cluster consisting of two NAS Platform 4080 server nodes managed by a SMU 400 on NAS Platform. The NAS Platform storage configuration is outlined below:

- 16 × 1.4 TB system drives consists of LUNs from the Hitachi Virtual Storage Platform G600 storage array.
- 22 TB storage pool consisting of 16 × 1.4 TB system drives.
- File system was created on NAS Platform 4080 named "HV3" to contain the virtual machine VHDX and configuration files for testing.
- File system was formatted with a 4 KB block size for performance.

Note — Set the Superflush option to 3 × 128 KB for each system drive to allow for better performance.

Server Blade Configuration

Two server blades on Hitachi Compute Blade 500 were used to host the test environment. The server blades had the Microsoft Hyper-V role installed. They were configured as a failover cluster. The server blades are listed in Table 6.

Table 6. Server Blade Configuration

Server Blade Type	Server Name	Role	Operating System
520H B2	CB500-11-B0	Hyper-V host server	Microsoft Windows Server 2012 R2 Datacenter
520H B3	CB500-11-B01	Hyper-V host server	Microsoft Windows Server 2012 R2 Datacenter

The virtual machines required for the infrastructure are listed in Table 7.

Table 7. Virtual Machine Configuration

Virtual Machine Name	Role	vCPU	Virtual Memory	Operating System
Hnas-cloud-DC1	Domain controller and DNS	1	4 GB	Microsoft Windows Server 2012 R2 Standard
Hnas-cloud-DC2	Domain controller and DNS	1	4 GB	Microsoft Windows Server 2012 R2 Standard
SCCOM	System Center Configuration Manager	2	8 GB	Microsoft Windows Server 2012 R2 Standard
SCVMM	System Center Virtual Machine Manager	2	8 GB	Microsoft Windows Server 2012 R2 Standard
SC-SQL	SQL Server for System Center databases	4	16 GB	Microsoft Windows Server 2012 R2 Standard

Each Microsoft Active Directory® domain controller was installed as a virtual machine on one of the hosts.

- Hnas-cloud-DC1 was installed on CB500-12-B0
- Hnas-cloud-DC2 was installed on CB500-12-B1

The domain controllers were not configured to fail over in the event of a server blade failure.

The other infrastructure virtual machines were installed on a Microsoft Hyper-V host. They were configured as roles in failover cluster manager to enable failover in the event of a blade server failure.

SAN Configuration

Two sets of zones were required for this environment.

- One set of zones was for the connections between the Hitachi Compute Blade 500 server blades and Hitachi Virtual Storage Platform G600.
- The other set of zones were for the connections between Hitachi NAS Platform 4060 heads and Virtual Storage Platform G600 storage system.

The zone configuration details are shown in Table 8.

Table 8. Fibre Channel Zone Configuration

Host	Host HBA or Port	Zone Name	Storage Ports
CB500-11-B0	HBA1_1	CB500_11_B0_HBA1_1_ASE42_26_3A	3A
	HBA1_2	CB500_11_B0_HBA1_2_ASE42_26_4A	4A
CB500-11-B1	HBA1_1	CB500_11_B1_HBA1_1_ASE42_26_5A	5A
	HBA1_2	CB500_11_B1_HBA1_2_ASE42_26_6A	6A
HNAS_522	1	HNAS_522_P1_ASE42_26_5A_6A	5A,6A
	2	HNAS_522_P2_ASE42_26_5A_6A	5A,6A
HNAS_523	1	HNAS_0523_P1_ASE42_26_5A_6A	5A,6A
	2	HNAS_0523_P2_ASE42_26_5A_6A	5A,6A

Network Configuration

To keep network configuration simple and reduce the chance of network configuration issues, all traffic used the same network interfaces and VLAN. This was acceptable for the test environment because only validation testing was being performed.

In a production environment or for performance testing, multiple network interfaces and VLANs should be used to separate the traffic. In a production environment, you must plan for the following types of traffic:

- **File share access —** Traffic between the Microsoft Hyper-V hosts and the file shares on Hitachi NAS Platform
- **Cluster communications —** Traffic between the nodes in the failover cluster
- **Management —** Traffic between the server or servers used to manage the environment and the devices being managed
- **Client-server communications —** Traffic between the clients and the Hyper-V host servers

Configuration Requirements

There are two configuration requirements to implement transparent failover for file shares hosted on the Hitachi NAS Platform 4080 cluster:

- Configure continuous availability on file share
- Configure Witness EVS

Transparent failover requires that the file shares used by the clients must have the continuous availability attribute enabled. Continuous availability enables files that are opened using SMB3 on a continuous available share to survive network or cluster node failures. If one cluster node fails, the client transparently reconnects to another node in the cluster without interruptions to the client applications.

For this testing, a file share with a CIFS name of HV3 was enabled for continuous availability.

To enable a file share for continuous availability, follow these steps in the GUI interface on NAS Platform 4080:

1. On the **Home** menu, click **File Services**, and then click **CIFS Share Details**.
2. To enable the share for continuous availability, select the **Ensure Share Continuously Available** check box, as shown in Figure 2.

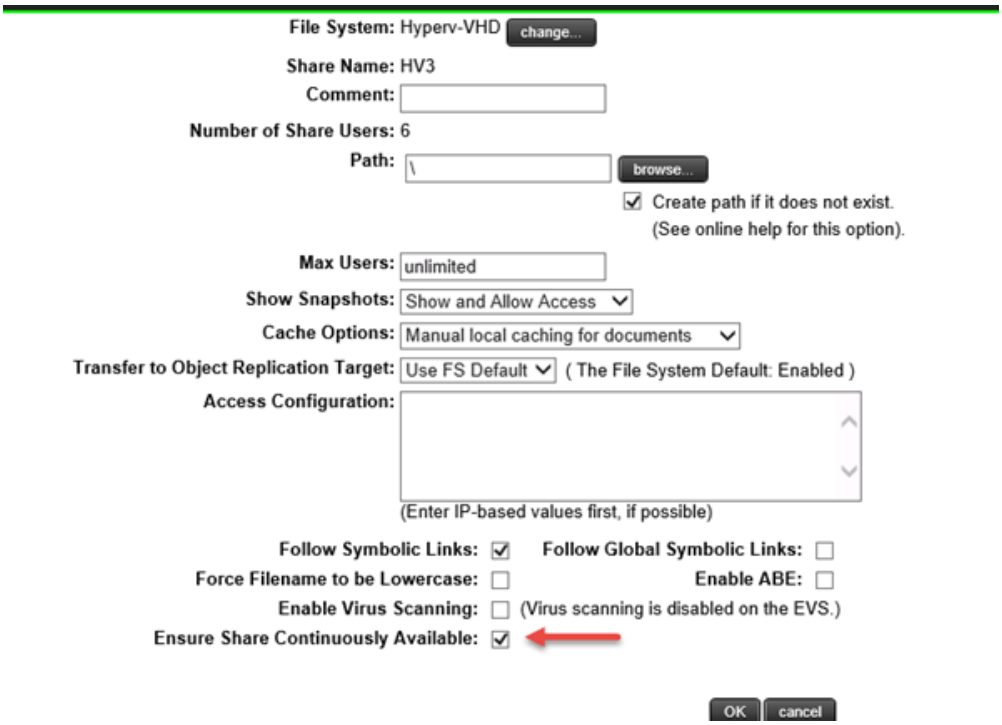

Figure 2

The witness server uses the service witness protocol to notify a client of any state changes on a continuously available server. The witness protocol enables the Hitachi NAS Platform node to notify the Microsoft Hyper-V server of any state changes without having to wait for the TCP connection to time out.

A witness EVS cannot be configured using the GUI. You must use the command line interface to configure a witness EVS to monitor a service EVS.

1. Create a witness EVS for the service EVS that contains the continuously available share by using the following command.

   ```
   evs create -l EVS01-WIT -i 172.27.4.173 -m 255.255.192.0 -p ag1 -w 1
   ```

   This command creates a witness EVS named EVS01-WIT with an IP address 172.27.4.123. It is a witness for EVS1, as specified in the **- w option**. This binds the witness EVS to service EVS 1.

2. Configure an ADS CIFS name for the witness EVS. Add the witness EVS by using the following command:

   ```
   cifs-name add -m ads -a <dc ip address> EVS01-WIT
   ```

Creating a witness EVS requires that, for proper operation, disabling DDNS for the EVS. This can be done using the GUI interface.

1. From the **Home** menu, click **File Services** and then click **CIFS Setup**.

2. To disable DDNS, click **Enable**, as shown in Figure 3. After you click **Enable**, it changes to **Disable**.

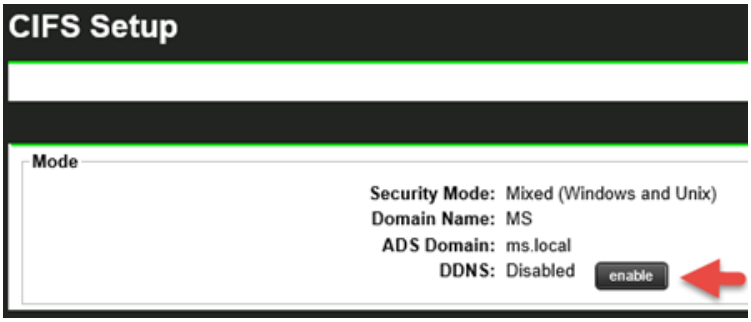

Figure 3

Test Methodology

Testing to validate SMB 3 transparent failover functionality on Hitachi NAS Platform in a Microsoft environment consisted of the following test cases:

1. Verify that the functions of transparent failover work when the EVS containing the file share for a Microsoft Hyper-V virtual machine **migrates** to another node in the Hitachi NAS Platform cluster while executing a heavy application workload.

   The expected result is that the workload keeps running without interruption.

2. Verify that the functions of transparent failover work when the node hosting the EVS containing the file share for a Microsoft Hyper-V virtual machine **reboots**, forcing a migration of the EVS to another node in the Hitachi NAS Platform cluster while executing a heavy application workload.

   The expected result is that the workload keeps running without interruption.

Workload Specifications

Standalone and multiple Web servers and SQL servers were deployed in two different I/O profiles.

To test the performance and resilience of transparent failover, Iometer, with an I/O profile of Microsoft SQL Server, generated Web server traffic to test transparent failover. This tool used an industry-standard Web server profile.

Results

With transparent failover enabled by making a file share continuously available and creating a witness EVS, the following are the testing results:

1. EVS containing the file share for a Hyper-V virtual machine was **migrated** to another node in the Hitachi NAS Platform cluster while a heavy application workload was executing.

   Result — Success. The workload continued to execute without any impact. No brownout or timeout was experienced by the workload

2. The node hosting the EVS containing the file share for a Micrsoft Hyper-V virtual machine was **rebooted** forcing a migration of the EVS to another node in the Hitachi NAS Platform cluster while a heavy application workload was executing.

   Result — Success. The workload continued to execute without any impact. No brownout or timeout was experienced by the workload

Analysis

Testing proved that using continuously available file shares with the Hitachi NAS Platform 4080 enables administrators to perform hardware or software maintenance on any cluster node without interrupting the server applications that are storing their data files on these file shares. Also, in case of a hardware or software failure, the server application nodes will transparently reconnect to another cluster node without interruption of the server applications.

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems Global Services website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate Resources website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services Education website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the Hitachi Data Systems website.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

AS-469-03. November 2016.