Storage Networking by Network Appliance and Cisco Systems:

High Availability for Network-Attached Storage

Manoj Joshi, Nicolas Breton, Jim Helfrich & Patrick Jones | Cisco Systems, Inc. and Network Appliance | TR-3115

Network Appliance Inc.

Table of Contents

1. Introduction

As data requirements increase at a rapid pace, end users demand reliable, high-performance, and global access to information from anywhere at anytime. IT managers constantly seek more affordable, manageable storage solutions that meet these user expectations. Network-attached storage (NAS) has led the way for the mainstream deployment of storage solutions that facilitate data consolidation and sharing. By leveraging well-understood technologies—Internet Protocol (IP), Gigabit Ethernet, Network File System (NFS), and Common Internet File System (CIFS)—NAS enables a flexible, robust storage solution that is easily managed and scaled, and contributes to a scalable and reliable network and storage infrastructure.

The IP/Gigabit Ethernet networking technology from Cisco Systems combined with Network Appliance™ storage appliances, delivers high performance and features such as scalability, availability, and security. The complementary solutions from Cisco and NetApp address the customer requirements for NAS in several dynamic application areas:

- **Internet E-Business Applications:** High-performance data sharing and scalable networked storage infrastructures for e-businesses
- **Business Applications in the Data Center:** Superior data availability and recoverability for enterprise business applications within a data center
- **Workgroup Collaboration:** High-performance data sharing across heterogeneous operating system environments
- **Distributed Storage over Secure WAN:** Collaboration among distributed sites with centralized administration and disaster recovery

This paper discusses the deployment of highly available Cisco and NetApp networked storage infrastructures for campus environments. Cisco and NetApp have jointly validated best design practices that integrate network and storage to enable the creation of highly available storage networking solutions. The solution combines the NetApp® NAS clustered architecture and the Cisco high-availability network architecture.

Based on these integrated architectures, two network design configurations were tested and evaluated. The configurations documented in this paper are based on the requirements for a typical large enterprise. The test results validate the combined Network Appliance and Cisco configurations for deploying highly available network storage solutions.

2. NAS Solution Architecture

The following storage networking architecture incorporates the Cisco network-layered architecture with the Network Appliance NAS architecture. The integrated solution enables reliable, high-performance, universal access to shared, consolidated information across a medium to large campus network.

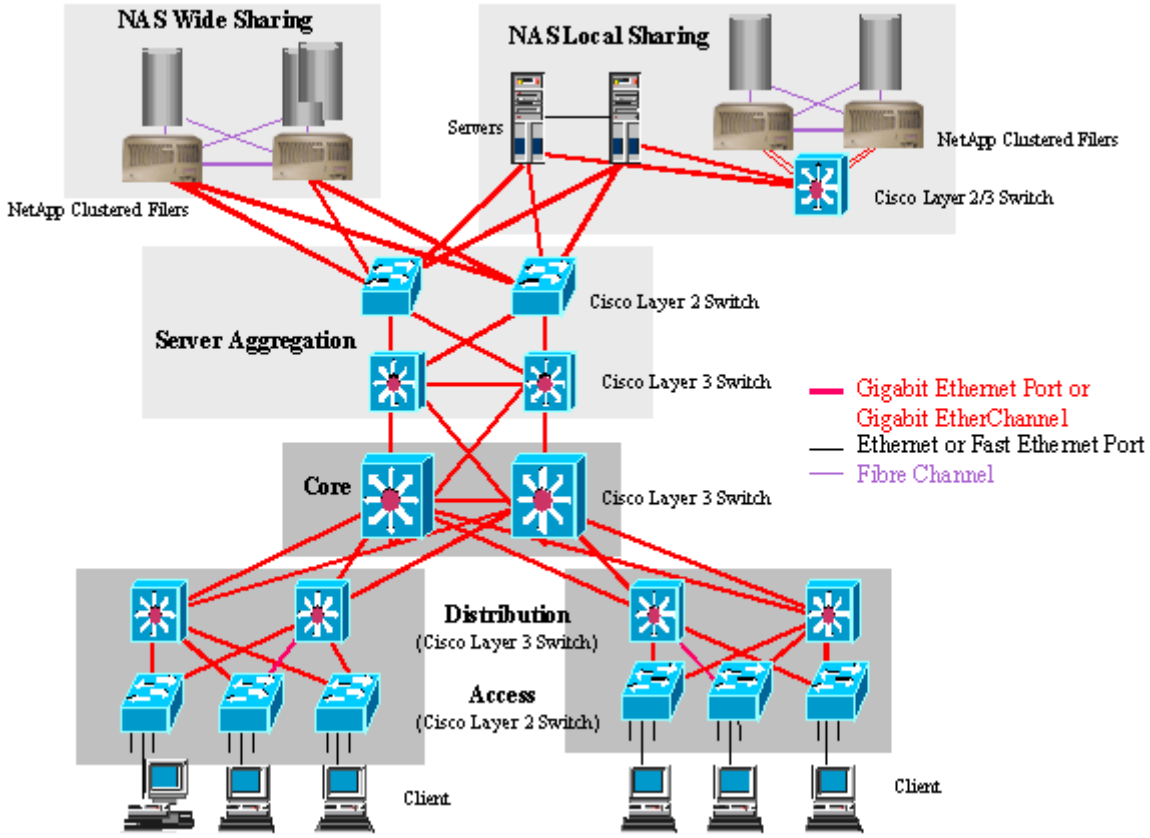

Figure 1 - Network-Attached Storage Solution Architecture

The combined Cisco and NetApp architecture provides the following benefits:

Availability: Redundant hardware components, in combination with multiple distributed devices, contribute to a highly available network infrastructure and data storage solution. The failure of any single hardware device or network link will not prevent access to data storage. Software features in both the network and storage layers provide for automatic failover and continuous data access.

Scalability: The modularity of the architecture allows for more accuracy in capacity planning at each layer. The solution is scalable at the different layers, permitting the enterprise to scale not only bandwidth but also the number of users and storage utilization. Load balancing can be carried out between different redundant devices and paths, efficiently managing traffic and optimizing link utilization.

Flexibility and Manageability: Simple but deterministic Layer 2 and Layer 3 paths make it easier to manage the network and deploy storage. Optional redundancy at each layer can be provided without breaking or disrupting the entire network. The flexibility of the architecture allows the addition of storage appliances to the network infrastructure based on the solution or application that the storage will support. Typical deployments include enterprise-wide sharing of data among numerous campus clients or local sharing between data center application servers and local storage appliances.

3. Cisco Systems Network Architecture

3.1. Access/Server Aggregation Layer

This layer has two functions: end-user connectivity such as workstations and printers (access) and connectivity for servers and network-attached storage (server aggregation). Scalability can be achieved at this layer with virtual LAN (VLAN) technology. VLAN technology provides good segmentation and management techniques for addressing. It permits effective control of the broadcast domains. Using Inter Switch Link (ISL) or IEEE 802.1Q VLAN trunking protocols, multiple VLANs can be carried over a single network path. The Cisco Spanning Tree Protocol (STP) implementations—Per VLAN Spanning Tree (PVST) and PVST+—allow redundant Layer 2 network paths between the server aggregation layer and the distribution layer. Cisco EtherChannel® technology allows network bandwidth to grow incrementally.

A VLAN can be spanned within the server aggregation layer. This allows dual-homed servers to be connected to the server aggregation layer in a more homogenous way, providing increased redundancy to the application servers at the network edge. Using Cisco EtherChannel technology, an aggregate bandwidth of up to 8Gbps can be scaled while connecting to servers and network-attached storage filers. EtherChannel technology also allows the load balancing of traffic based on MAC address, IP source/destination address, and Layer 4 port numbers.

3.2. Distribution Layer

The distribution layer acts as a clear demarcation between the core access layer. It aggregates Layer 2 traffic from multiple access layer switches, terminates VLANs and subnets, and directs traffic as needed to the core.

The distribution layer provides the mechanism for first-hop redundancy to the clients/servers by configuring Hot Standby Routing Protocol (HSRP). Additionally, HSRP optional features such as "interface tracking" and "preempt delay" ensure increased availability. The network can quickly converge in the event of a connectivity failure to the core layer, and allow reconvergence of Layer 3 due to bringing up failed distribution switches.

In the event of a failure in the distribution layer, the overall convergence time can be minimized by overlaying Layer 2 (Spanning Tree Root [Primary]) and Layer 3 (HSRP [Primary]) on the same distribution switch. Scalable routing protocols such as Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP) can be fine-tuned to achieve the best possible convergence time and to optimize load balancing between the distribution and core layers.

3.3. Core Layer

Typically, the core layer is a Layer 3 high-speed backbone for the campus network with low latency and high packet throughput. It acts as an access point for intranet connectivity and aggregates traffic from multiple distribution layers.

4. Network Appliance Storage Architecture

4.1. Availability: Filer Cluster

Two NetApp filers can be configured as a cluster to provide increased protection against hardware failures. Clustered NetApp filers are connected through an interconnect adapter and cables, and are configured so that both filers share access to the same set of Fibre Channel disks and networks. Each filer uses the cluster interconnect to continually monitor the availability of the

partner filer. The interconnect is also used to mirror each filer's NVRAM log data and to synchronize the time of the clustered partners. Fibre Channel loops connect each filer to its own disks, and a separate Fibre Channel loop provides a connection to its partner's disks (see Figure-1).

Each filer has primary responsibility for a subset of the disks and both can operate independently. The NetApp cluster architecture is an active/active configuration. During normal operation, both filers are operating and serving data from their individual disk arrays. If a system failure occurs on one filer in a cluster, the partner filer will perform a *takeover* of the failed filer functions and provide client access to the data on the failed filer's disk arrays. The partner filer maintains its own network identity and its own primary functions, but also assumes the network identity of the failed machine and handles the added functionality through a virtual filer. Redundant disk, fan, or power supply failures are handled independently, in the same manner as with a standalone filer; these failures do not trigger failovers. In addition to an automatic takeover, a manual takeover can be forced at any time. This allows some scheduled filer maintenance tasks to be completed without interrupting data services.

NetApp clustered partners protect against a failure of a filer system unit, not a single network interface failure. Network connection problems are better handled at the interface level. Data ONTAP™ software allows filer network interface cards to be configured in multihomed configurations with an active interface that fails over to a standby interface whenever a loss of network link occurs. Multiple interfaces can be bundled to form a compatible Cisco EtherChannel. In this configuration, all interfaces are active and the switch controls the link failover and load balancing.

4.2. Scalability: Network Storage

Network Appliance storage provides a very flexible configuration for building a scalable storage network. NetApp filers connect directly to the network and are accessed using the industry-standard network file system protocols: NFS and CIFS. By separating the data storage and file system operations from clients and application servers, NetApp filers allow the network storage farm to transparently scale in capacity and file system processing power.

When more storage is needed, additional disks can be added to one or multiple NetApp filers. There is no need to reconfigure application server hardware or pre-allocate storage for a specific client operating system or file system type. Data is immediately available to all network clients and servers. In most cases, this storage can be added on-the-fly without disrupting applications. A single filer can scale to multiple terabytes. Multiple filers can be used to scale storage into very large storage farms.

A NetApp filer provides a complete file store and handles all the necessary processing and management of data. File system operations, RAID protection, quota management, and multiprotocol access are all handled by the storage appliance. Application servers and clients can devote their resources to running applications. When additional application servers are needed, they can be transparently added. If more file system operations and storage processing power is needed, additional filer processor units can be added to the storage farm. A storage farm architecture can be infinitely scaled for most applications and networks.

4.3. Flexibility: Wide Sharing and Local Sharing Deployments

The flexibility of a storage network allows for deployments in a variety of different applications and network configurations. Enterprise, or wide area, sharing enables information stored on a filer to be accessible from numerous locations and clients. Different types of clients and applications can all share data stored in home directories and project directories. The client machines

communicate directly with the filer. The client desktop machines can be in the same building or on opposite sides of the world from the filer (see Figure-1).

The high reliability and data management features of NetApp filers make them suitable storage solutions for mission critical server applications such as database, e-mail, and Web servers. In a local sharing environment, the filer and application server are located in the same data center and communicate on a dedicated local server farm network. Clients accessing the application server use separate wide-sharing network connections to send their requests to the application server. All filer-server communications are carried out by the application servers over the dedicated storage network. The dedicated network connections between the application servers and the filers provide high reliability and high performance.

5. Integration and Testing

The previously described Cisco network-layered architecture is essential to the deployment of an enterprise-wide high-availability NAS solution. Network Appliance and Cisco have jointly validated two high-availability network and storage designs. Both make use of clustered NetApp filers and Cisco Catalyst® Series switches.

The integration and testing of the high-availability network and storage solutions requires that failure scenarios of each of the redundant network and storage building-block elements be validated. Tests cover the failures of filer cluster units, network links, and Catalyst series switches. A network *ping* test, from the Microsoft® Windows® 2000 and UNIX® clients, was utilized to determine the average failover and recovery times. Failover time was measured for each test scenario that involved removing an element from the test network. Recovery time was measured when the failed element was later reinserted into the test network. An NFS file copy was also used to test data access continuity during each failure and recovery scenario. The test results show how the network and storage designs react to various network link and device failures, and validate the interoperability of Cisco and Network Appliance solutions.

Two design approaches to building a highly available network infrastructure were tested. These two approaches include a network design with multiple distributed switches and another configuration with redundant hardware in a single Catalyst switch. Each topology applies to a different user scenario and addresses different levels of high-availability requirements. The two tested configurations are referred to as:

- Distributed server aggregation layer (Network Design Option 1)
- Redundant server aggregation layer (Network Design Option 2)

5.1. Network Design Option 1—Distributed Server Aggregation Layer

The design shown in Figure 2 uses a highly redundant Catalyst 6500 switch in the distribution layer and redundant Catalyst 4006 switches in the access/server aggregation layer. The distribution and access/server aggregation layers are connected with Gigabit Ethernet.

Each dual-homed clustered filer is connected to the two Catalyst 4006 switches at the access/server aggregation layer. Both Catalyst 4006 switches are configured with Spanning Tree Protocol enabled, the "PortFast" feature enabled for the ports connecting to filers, and "UpLinkFast" feature enabled for the uplink ports connecting to the Catalyst 6500 switches. Also, both Catalyst 6500 switches have the Spanning Tree Protocol and "BackBoneFast" feature enabled. To overlap Layers 2 and 3 for fast and predictable network convergence, one of the Catalyst 6500 switches is configured with Spanning Tree Root [Primary] and HSRP [Active].

The Network Appliance dual-homing method makes use of Layer 2 connectivity. In this configuration, only one interface is active and the second is in standby mode. In the event that an active interface/link fails, the standby interface takes over the active role with minimum network interruption. The filer's Data ONTAP software monitors the active link and handles the failover. The dual-homed filer configuration has been tested and validated with both Fast Ethernet and Gigabit Ethernet connectivity.

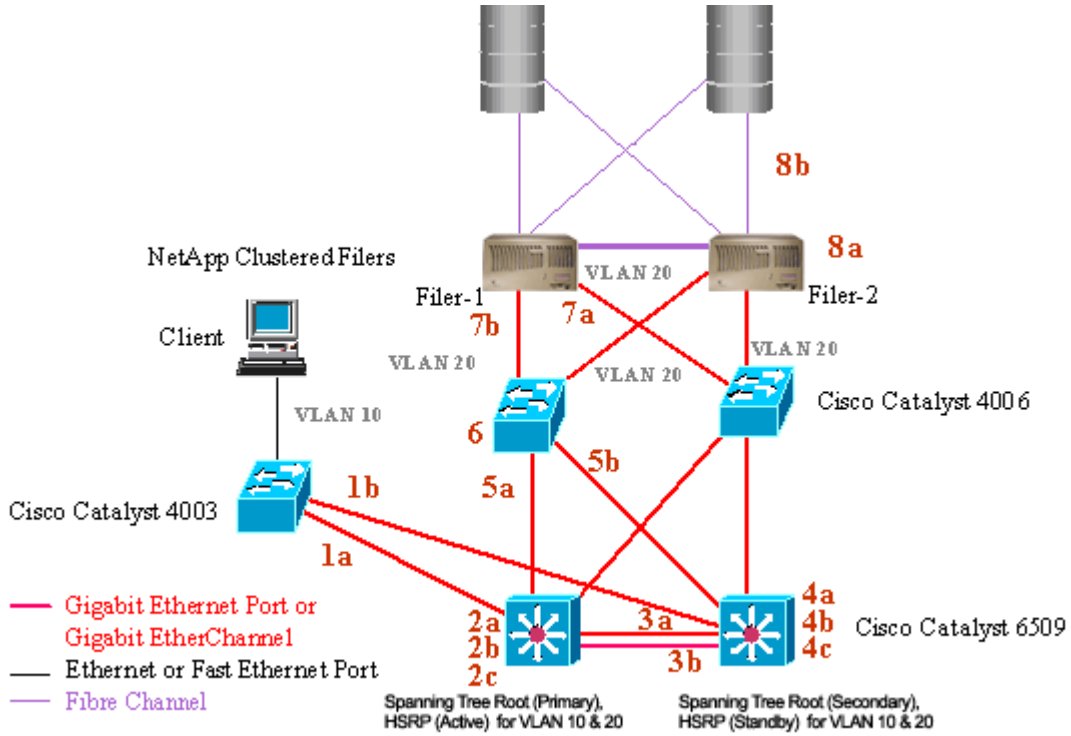

Figure 2 - NAS Network Design Option 1

5.2. Test Results for Network Design Option 1

The results indicate the time taken for the network to converge after a failure, and for recovery after reinserting the failed link or device. The network convergence time is measured from the Windows 2000 client, with the client continuously *pinging* both NetApp filer cluster partners at one-second intervals. Different failure scenarios were simulated, failing one element or segment during each test. During all test scenarios, a series of NFS file copy operations was performed to validate that continuous data access was maintained during failure and recovery phases. Network Design Option 1 passed all failure and recovery tests.

The table below details the 17 test scenarios. The table indexes 1a through 8b, which correspond to the physical locations depicted in Figure 2. The tabular results apply to the test configurations only. Variations in equipment and configurations will have an effect on the failover times.

	Type of Failure and Recovery	Convergence Time (in Sec)	Feature Responsible for Convergence
1a	Fail the "forwarding" uplink between access and distribution (Layers 2 and 3 active) switch	3	Spanning Tree; UpLinkFast
	Restore the failed uplink between access and distribution (Layers 2 and 3 active) switch	0	Spanning Tree

1b	Fail the "blocking" uplink between access and distribution (Layers 2 and 3 standby) switch	0	Spanning Tree
	Restore the failed uplink between access and distribution (Layers 2 and 3 standby) switch	0	Spanning Tree
2a	Fail the active supervisor on the distribution (Layers 2 and 3 root) switch	1	HSRP
	Restore the failed supervisor on the distribution (Layers 2 and 3 root) switch	0	Supervisor HA Protocol
2b	Fail the router on the distribution (Layers 2 and 3 root) switch	1	HSRP
	Restore the failed router on the distribution (Layer 2 root) switch	1	HSRP
2c	Fail the distribution (Layers 2 and 3 root) switch	1	HSRP
	Restore the failed distribution switch	8	Spanning Tree; HSRP
3a	Fail one of the links in the EtherChannel between distribution switch	0-1	EtherChannel
	Restore the failed link in the EtherChannel between distribution switch	0-1	EtherChannel
3b	Fail both links in the EtherChannel between distribution switch	4	Spanning Tree; BackBoneFast
	Restore both the failed links in the EtherChannel between distribution switch	0	Spanning Tree
4a	Fail the active supervisor on the distribution (Layers 2 and 3 standby) switch	0	Supervisor HA Protocol
	Restore the failed supervisor on the distribution (Layers 2 and 3 standby) switch	0	Supervisor HA Protocol
4b	Fail the router on the distribution (Layers 2 and 3 standby) switch	0	HSRP
	Restore the failed router on the distribution (Layers 2 and 3 standby) switch	0	HSRP
4c	Fail the distribution (Layers 2 and 3 standby) switch	0	Spanning Tree; HSRP
	Restore the failed distribution (Layers 2 and 3 standby) switch	0	Spanning Tree; HSRP
5a	Fail the "forwarding" uplink between server aggregation and distribution (Layers 2 and 3 active) switch	3	Spanning Tree; UpLinkFast
	Restore the failed uplink between server aggregation and distribution (Layers 2 and 3 active) switch	0	Spanning Tree
5b	Fail the "blocking" uplink between server aggregation and distribution (Layers 2 and 3 standby) switch	0	Spanning Tree
	Restore the failed uplink between server aggregation and distribution (Layers 2 and 3 standby) switch	0	Spanning Tree
6	Fail the server aggregation switch	3	NetApp Layer 2 [Single-Mode VIF]

	Restore the failed server aggregation switch	0	Spanning Tree; PortFast
7a	Disconnect the standby link between filer and server aggregation switch	0	NetApp Layer 2 Protocol [Single -Mode VIF]
	Restore the disconnected link between filer and server aggregation switch	0	Spanning Tree; PortFast
7b	Disconnect the active link between filer and server aggregation switch	1-2	NetApp Layer 2 Protocol [Single -Mode VIF]
	Restore the disconnected link between filer and server aggregation switch	0	Spanning Tree; PortFast
8a	Fail the filer head [switch off the filer]	45	NetApp Cluster Monitor
	Restore the filer head [switch "on" the filer]	75	NetApp Cluster Monitor
8b	Disconnect the link between filer and disk array	120-180	NetApp Cluster Monitor
	Reconnect the link between filer and disk array	0	NetApp Cluster Monitor

5.3. Network Design Option 2—Redundant Server Aggregation Layer

The design shown in Figure 3 makes use of highly redundant Catalyst 6500 switches in the distribution and server aggregation layers, with redundant Catalyst 4006 switches in the access layer. The distribution and access/server aggregation layers are connected with Gigabit Ethernet.

The NetApp filer cluster connects to highly redundant Catalyst 6500 switches with dual supervisor (high-availability option enabled) through Gigabit EtherChannel. The Gigabit Ethernet ports connecting to the filers are bundled across separate blades using Cisco EtherChannel, providing physical redundancy and load balancing. The uplink ports connecting to Catalyst 6500 in the distribution layer are configured with the *UpLinkFast* feature enabled, and the ports connecting to filers are configured with the *PortFast* option enabled for fast convergence in the event of a link failure.

To take advantage of redundancy and load balancing, all switches have Spanning Tree Protocol enabled. The Catalyst 6500 switches in the distribution layer are configured with the *BackBoneFast* feature, which provides fast Layer 2 convergence due to indirect failure. To overlap Layers 2 and 3 for fast and predictable network convergence, one of the Catalyst 6500 switches in the distribution layer is configured with Spanning Tree Root [Primary] and HSRP [Active].

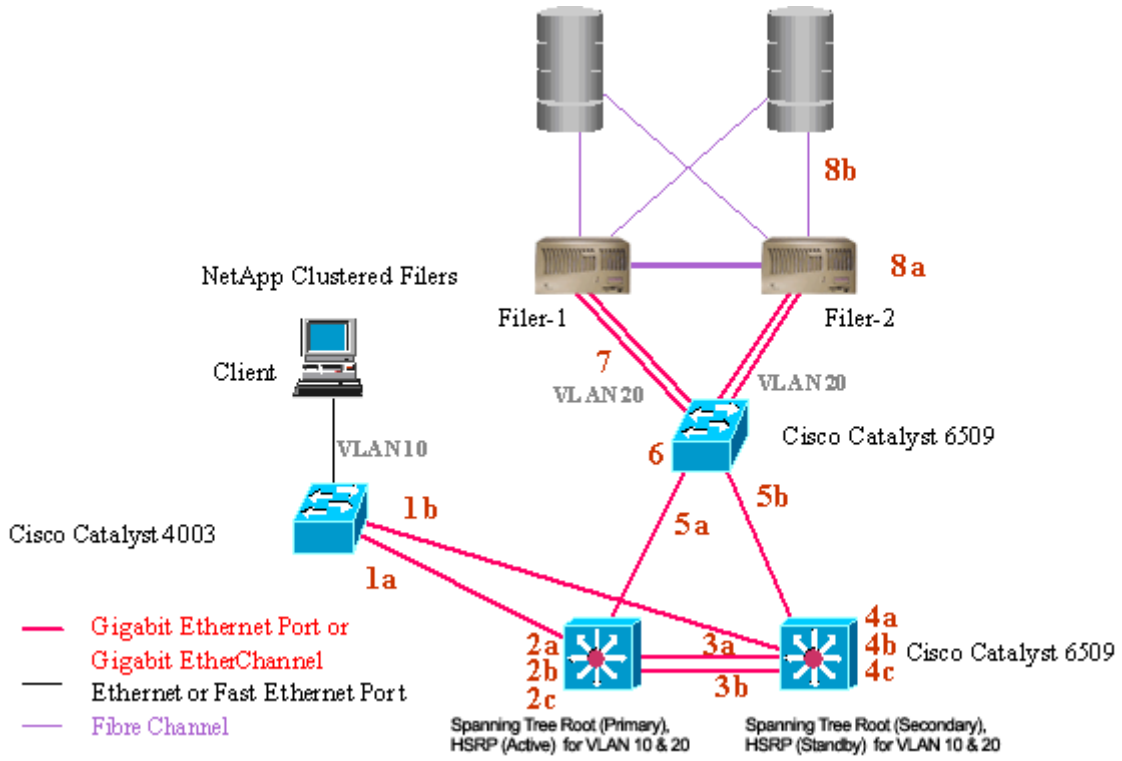

Figure 3 - NAS Network Design Option 2

5.4. Test Results for Network Design Option 2

The results indicate the time taken for the network to converge after a failure, and for recovery after the failed link or device is reinserted. The network convergence time is measured from the Windows 2000 client, with the client continuously *pinging* both NetApp filer cluster partners at one-second intervals. Different failure scenarios were simulated, failing one element or segment during each test. During all test scenarios, a series of NFS file copy operations was performed to validate that continuous data access was maintained during failure and recovery phases. Network Design Option 2 passed all failure and recovery tests.

The table below details the 17 test scenarios. The table indexes 1a through 8b, which correspond to the physical locations depicted in Figure 3. The tabular results apply to the test configurations only. Variations in equipment and configurations will have an effect on the failover times.

	Type of Failure and Recovery	Convergence Time (in Sec)	Feature Responsible for Convergence
1a	Fail the "forwarding" uplink between access and distribution (Layers 2 and 3 active) switch	3	Spanning Tree; UpLinkFast
	Restore the failed uplink between access and distribution (Layers 2 and 3 active) switch	0	Spanning Tree
1b	Fail the "blocking" uplink between access and distribution (Layers 2 and 3 standby) switch	0	Spanning Tree
	Restore the failed uplink between access and distribution (Layers 2 and 3 standby) switch	0	Spanning Tree
2a	Fail the active supervisor on the distribution (Layers	1	HSRP

	2 and 3 root) switch		
	Restore the failed supervisor on the distribution (Layers 2 and 3 root) switch	0	Supervisor HA Protocol
2b	Fail the router on the distribution (Layers 2 and 3 active) switch	1	HSRP
	Restore the failed router on the distribution (Layer 2 active) switch	1	HSRP
2c	Fail the distribution (Layers 2 and 3 active) switch	1	HSRP
	Restore the failed distribution (Layers 2 and 3 active) switch	8	Spanning Tree; HSRP
3a	Fail one of the links in the EtherChannel between distribution switch	0-1	EtherChannel
	Restore the failed link in the EtherChannel between distribution switch	0-1	EtherChannel
3b	Fail both links in the EtherChannel between distribution switch	4	Spanning Tree; BackBoneFast
	Restore both the failed links in the EtherChannel between distribution switch	0	Spanning Tree
4a	Fail the active supervisor on the distribution (Layers 2 and 3 standby) switch	0	Supervisor HA Protocol
	Restore the failed supervisor on the distribution (Layers 2 and 3 standby) switch	0	Supervisor HA Protocol
4b	Fail the router on the distribution (Layers 2 and 3 standby) switch	0	HSRP
	Restore the failed router on the distribution (Layers 2 and 3 standby) switch	0	HSRP
4c	Fail the distribution (Layers 2 and 3 standby) switch	0	Spanning Tree; HSRP
	Restore the failed distribution (Layers 2 and 3 standby) switch	0	Spanning Tree; HSRP
5a	Fail the "forwarding" uplink between server aggregation and distribution (Layers 2 and 3 active) switch	3	Spanning Tree; UpLinkFast
	Restore the failed uplink between server aggregation and distribution (Layers 2 and 3 active) switch	0	Spanning Tree
5b	Fail the "blocking" uplink between server aggregation and distribution (Layers 2 and 3 standby) switch	0	Spanning Tree
	Restore the failed uplink between server aggregation and distribution (Layers 2 and 3 standby) switch	0	Spanning Tree
6	Fail the active supervisor on the server aggregation switch	2-3	Supervisor HA Protocol
	Restore the failed supervisor on the server aggregation switch	0	Supervisor HA Protocol
7	Fail one of the links in the EtherChannel between filer and server aggregation switch	0-4	EtherChannel

	Restore the failed link in the EtherChannel between filer and server aggregation switch	0	EtherChannel
8a	Fail the filer head [switch off the filer]	45	NetApp Cluster Monitor
	Restore the filer head [switch "on" the filer]	75	NetApp Cluster Monitor
8b	Disconnect the link between filer and disk array	120-180	NetApp Cluster Monitor
	Reconnect the link between filer and disk array	0	NetApp Cluster Monitor

6. Conclusion

Continuous data access to a network storage solution involves two main elements: a redundant and fault-tolerant network infrastructure, and a highly available storage solution with built-in data protection features. The testing outlined in this paper focused on integrating the two elements into a high availability network storage solution with no single point of failure. The solutions integrate Cisco network designs for campus environments with Network Appliance storage clusters. The published test results validate the effectiveness of a Cisco and NetApp high-availability data access and storage solution.

7. Appendix A: Cisco Systems Products

7.1. Catalyst 6000 Family

The Catalyst® 6000 family, which consists of the Catalyst 6000 and 6500 series, delivers high-performance, multilayer switching solutions for enterprise and service provider networks. Designed to address the increased requirements for gigabit scalability, high-availability, and multilayer switching in backbone/distribution and access/server aggregation environments, the Catalyst 6000 family provides exceptional scalability and price/performance ratios. These switches support a wide range of interface densities, performance, and high-availability options. Combining superior control-plane and packet-forwarding scalability with a rich set of intelligent services, the Catalyst 6000 family delivers the foundation for New World enterprise and service provider solutions such as converged voice, video, data, and storage services.

The Catalyst 6000 family now enables service provider and enterprise environments to take the next step in scalable network architectures. Combining the benefits of scalable performance, scalable control plane, and extensive intelligent network services, the Catalyst 6000 family provides the framework for next-generation network build outs and mission-critical applications that require 24x7 availability. The Switch Fabric Module for the Catalyst 6500 series switches delivers best-in-class bandwidth performance for today's most advanced networks, with switching capacity that scales to 256Gbps. Catalyst 6500 series high-performance Gigabit Ethernet switching modules with local or distributed forwarding are ideal for deployment in gigabit backbone and server farm configurations or for the aggregation of high-density 10/100Mbps wiring closets.

The Catalyst 6500 family provides various high-availability options, including redundant supervisor, power supply, and fan modules. The redundant supervisor option facilitates change management for improved network serviceability.

Industry-leading Cisco Express Forwarding (CEF) and Gigabit Ethernet modules provide scalable control-plane, packet-forwarding performance and intelligent services for security and quality of service (QoS) to deliver next-generation solutions for dynamic service provider and enterprise networks.

Cisco EtherChannel technology allows Fast Ethernet and Gigabit Ethernet port aggregation up to 8Gbps and provides load balancing based upon source and destination IP and MAC addresses and Layer 4 port numbers.

More information about the Catalyst 6000 families is available at www.cisco.com. For more detail on high-availability network designs, please refer to the Cisco white paper on gigabit campus design at http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/camp_wp.htm.

7.2. Catalyst 4000 Series

The Catalyst 4000 series provides an advanced enterprise switching solution for wiring closets and data centers. Intelligent services can be delivered using the 24Gbps bandwidth architecture for 10/100/1000Mbps Ethernet switching. The Catalyst 4000 series Ethernet switches offer a broad range of port density and functionality.

The Catalyst 4000 series integrates the features and performance required in today's wiring closets, while ensuring the seamless introduction of tomorrow's enterprise enhancements. Wiring closet requirements include support for wired and wireless PC connections, IP telephony appliances and gateways, multicast streaming applications (such as Cisco IP/TV®), and gigabit to the desktop. The Catalyst 4000 modular chassis solution offers the flexibility to meet these needs now and in the future.

More information about the Catalyst 4000 series is available at www.cisco.com.

8. Appendix B: Network Appliance Products

8.1. F700 and F800 Series Storage Appliances

The Network Appliance F700 series and F800 series of filer storage appliances are the building blocks that allow companies to simplify, share, and scale data storage. NetApp filers serve as high-availability solutions for consolidating data and simplifying data management. The appliances are easy to manage and smoothly integrate into today's network-centric data center architectures. Using industry-standard file system access protocols, NetApp filers deliver very high-performance data access. Multiprotocol heterogeneous data sharing enables simultaneous file system access for UNIX, Windows NT®, and Web-based servers and clients. Built-in NetApp Data ONTAP software also offers data management tools and data protection solutions that enhance high availability for continuous data access.

NetApp filers reduce costly downtime and maximize access to mission-critical data. Built-in RAID protects against data loss and downtime due to disk failures. In the event of a failure, automatic reconstruction takes place on a hot spare disk with notification sent to the system administrator. Replacement disks are hot swappable. Disk scrubbing ensures data integrity and battery-backed NVRAM provides additional data protection. Redundant hardware components are built into each filer unit. Filers can also be configured in clustered failover pairs for a higher level of hardware redundancy and data availability.

Easy to manage and maintain, NetApp filers can be administered using only a few simple commands. System software can be updated in less than three minutes, and system administration can be handled using a Web browser. Redundant, hot-pluggable power supplies and cooling fans allow for online parts replacement. Retractable system trays provide easy access to internal components and decrease repair and service time. An auto support e-mail feature automatically reports potential problems to Network Appliance technical support and local IT administrators. NetApp filers can also be monitored and integrated with Simple Network Management Protocol (SNMP) tools.

The Network Appliance Data ONTAP software supports the following features for enhanced availability and delivering enterprise-class solutions:

- **Snapshot™ technology:** enables near-instantaneous, transparent online backup, storing up to 255 read-only versions of each data volume. End users can quickly recover deleted or modified files without administrative assistance or restore from tape backup. The Snapshot function requires minimal disk space and causes no disruption of service.

- **SnapRestore® software:** allows any system to revert back to a specified data volume Snapshot for instant file-system recovery. Terabytes can be recovered in minutes rather than hours, without going to tape.

- **SnapMirror® software:** provides remote mirroring at high speeds over a LAN or WAN. The asynchronous mirroring can be used for disaster recovery, replication, backup, or testing on a nonproduction system.

Additional information about Network Appliance data storage and management products is available at www.netapp.com.

9. Appendix C: Glossary

Common Internet File System (CIFS): Microsoft's file-sharing protocol that evolved from SMB.

Network-Attached Storage (NAS): A storage device, commonly referred to as a filer, that is connected to a network.

NetApp: The trademarked short form of the company name Network Appliance, Inc.

Network File System (NFS): A protocol for networking computers in a UNIX environment.

Network Interface Card (NIC): A printed circuit board that connects a computer or other node to a network. Also known as a "network adapter."

Non-volatile Random Access Memory (NVRAM): A type of computer memory that retains data in the event of a loss of power. In NetApp filers, NVRAM is used for logging incoming write data and requests.

Redundant Array of Independent Disks (RAID): NetApp appliances use RAID 4, which protects against disk failure by computing parity information based on the contents of all the disks in the array.

Simple Network Management Protocol (SNMP): A standard Internet protocol that facilitates communications between a system being managed and the management console or framework.

For Cisco glossary of terms, please refer to: http://www.cisco.com/warp/public/5/glossaries/logos/