An Open Standards Approach to Network-Centric Storage

Iftikhar Ahmed, Rajesh Godbole & Shiva Vishwanathan | Cisco Systems, Inc. and Network Appliance | TR-3121

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Network Appliance Inc.

Table of Contents

1. Executive Summary

As enterprises experience the widespread adoption of Internet business applications, the explosion in data reflects the continual proliferation and dependence on data storage infrastructure. Driving this growth is the increasing deployment of solutions such as e-commerce, customer care, workforce automation, and e-learning. This rapid growth, along with the escalating management costs associated with the storage infrastructure, has resulted in significant interest in moving from a direct-attached storage model to a more scalable and manageable networked storage model.

IDC predicts[1] that networked storage is expected to grow from US$6 billion in 1999 to approximately US$32 billion in 2004 (CAGR = 51.3%) while direct-attached storage is expected to shrink from US$14 billion in 1999 to approximately US$9 billion in 2004 (CAGR = -10%).

Storage networking enables companies to extend their storage networks beyond isolated islands in the data center to campus, metropolitan, and wide area environments. Storage networking has emerged as an increasingly strategic component of the IT infrastructure, addressing the following needs:

- Improve data availability and integrity
- Scale, share, and optimize storage resources
- Simplify storage management
- Minimize total cost of ownership (TCO) for storage

The goal of a consolidated storage network is to provide a framework for uniting multiple storage architectures-including direct-attached storage, network-attached storage (NAS), and Fibre Channel storage area networks (SANs)-into a single, well-managed, scalable, and extensible storage infrastructure. The Cisco Systems storage networking initiative is a comprehensive set of technologies, products, and partner solutions that enables the development of a networked storage infrastructure based on open architectures and industry standards.

This paper, which is jointly produced by Cisco Systems, Inc., and Network Appliance, Inc., (NetApp) highlights key market drivers for IP-connected network-attached storage, and elaborates on some common deployment scenarios in customer environments. Cisco and NetApp share a common vision to develop joint solutions based on industry standards that enable universal access to storage in both distributed and centralized storage environments.

NAS has led the way for the mainstream deployment of IP-based storage consolidation and file sharing. Using well-understood technologies, such as IP, Gigabit Ethernet, Network File System (NFS), and Common Internet File System (CIFS), NAS provides a flexible storage solution that is easily scaled and managed across large enterprise environments. NetApp offers NAS appliances that deliver the high performance and availability required for business-critical applications.

Coupled with Gigabit Ethernet, IP multilayer switching, and routing platforms from Cisco, the combined solutions leverage the already deployed and trusted IP infrastructure. This helps deliver an exceptional low TCO by reducing associated operational costs, complexity, and deployment time.

The Cisco and NetApp joint solutions address several customer requirements for network-attached storage:

- **Internet e-business applications**—High-performance data sharing and scalable networked storage infrastructure 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

2. The Need for Storage Networking

Globalization of businesses as well as the nature of e-business initiatives demand around-the-clock, 24x7 operation. Any downtime results not just in productivity loss but also significant revenue impact. Therefore, data availability has emerged as a critical corporate requirement for e-business IT infrastructure.

To address escalating storage demands, most corporations store enterprise data captive in servers or within a storage subsystem that is directly attached to a server. As the volume of storage data increases, this server-centric architecture has proven difficult to scale, manage, and deliver 24x7 availability.

Networked storage promises to reduce the cost and complexity associated with delivering highly available and scalable storage services. This network-centric model, which is termed storage networking, can be described as the software and hardware that enables storage to be consolidated, shared, accessed, and managed over a networked infrastructure.

Storage networking focuses on accelerating the convergence of storage and networks based on open architecture and industry standards. NAS is an implementation model of storage networking that delivers convergence of storage with IP-based networks.

3. Cisco and Network Appliance: An Integrated Approach

Successful NAS deployments hinge upon the integration of storage appliances with a scalable and intelligent network infrastructure. Cisco and Network Appliance are collaborating to develop, deploy, and support networked storage solutions that provide high-performance data access, data protection, and disaster recovery for shared storage resources over an IP network.

Network Appliance™ NAS solutions provide high-performance data sharing and administration for workgroup, departmental, and enterprise environments. These solutions leverage open standards such as NFS, CIFS, and the Network Data Management Protocol (NDMP) and deliver advanced capabilities such as clustering and remote replication. Using an IP network infrastructure to provide universal access and interconnections for NAS allows storage access, replication, and backup over enterprisewide networks. The joint Cisco and Network Appliance solutions provide the following benefits to customers:

- **Scalable storage infrastructure**—Cisco proven network offerings such as Ethernet switching in the LAN and routing across enterprise networks, and NetApp® storage appliances, deliver a networked storage architecture that is highly scalable not only within the data center but also across campus and wide area networks. In addition, this architecture allows storage and server farms to scale independently of each other.

- **Improved business continuance and data protection**—A redundant or clustered networked storage architecture ensures high availability of data. A consolidated storage network, combined with software capabilities, enables global management and remote backup/restore and disaster recovery implementation for NAS environments.

- **Consolidated network design, administration, and support**—Customers can leverage their existing IP network architectures and expertise instead of building separate storage networks based on technologies that need a separate administration and management framework.

- **Lower cost of ownership**—Consolidation, sharing, and access of storage over a storage network reduces IT infrastructure management costs, while using storage resources more efficiently. The simple-to-use appliance approach of NetApp and commonly available expertise on Cisco end-to-end network offerings make these joint solutions easy to manage and administer, thus offering the lowest possible TCO.

- **Intelligent IP services**—Various intelligent services at the IP layer ensure continuous data availability, protection, and scalability for the solution architectures. Bandwidth can be scaled while leveraging intelligent network services such as Layer 3 switching, quality of service (QoS), caching, server load balancing, and security.

4. Deployment Scenarios

The following sections describe typical customer deployments for a combined networked storage solution from Cisco and Network Appliance including solution architecture, components, and design considerations:

- **Internet and E-business Applications**

- **Business Applications in the Data Center**

- **Workgroup Collaboration**

- **Distributed Storage over Secure WAN**

These deployment examples can be used as a baseline reference and can be tailored to suit specific customer environments.

4.1. Internet and E-business Applications

E-business applications require highly scalable and available networked storage infrastructures that deliver exceptional performance for the Web, application, and database tiers. Examples include:

- An Internet portal that delivers Web-based applications such as e-mail, online calendar, and personalized multimedia content for a user community

- A business-to-business (B2B) application that connects suppliers and customers to an enterprise

- Intranet applications, such as workforce automation, for enhanced productivity

Solution Description
The solution architecture (Figure 4.1) illustrates a typical Internet infrastructure deployment in an e-business environment. It highlights the networked storage components along with some of their associated design considerations.

Delivers an efficient architecture for the Web tier.
In a tiered architecture, as shown in the diagram, the options to deploy storage at the Web tier include deployment of Web servers with either of the following:

- Direct-attached storage to store all Web content
- Consolidated networked storage connected to these Web servers

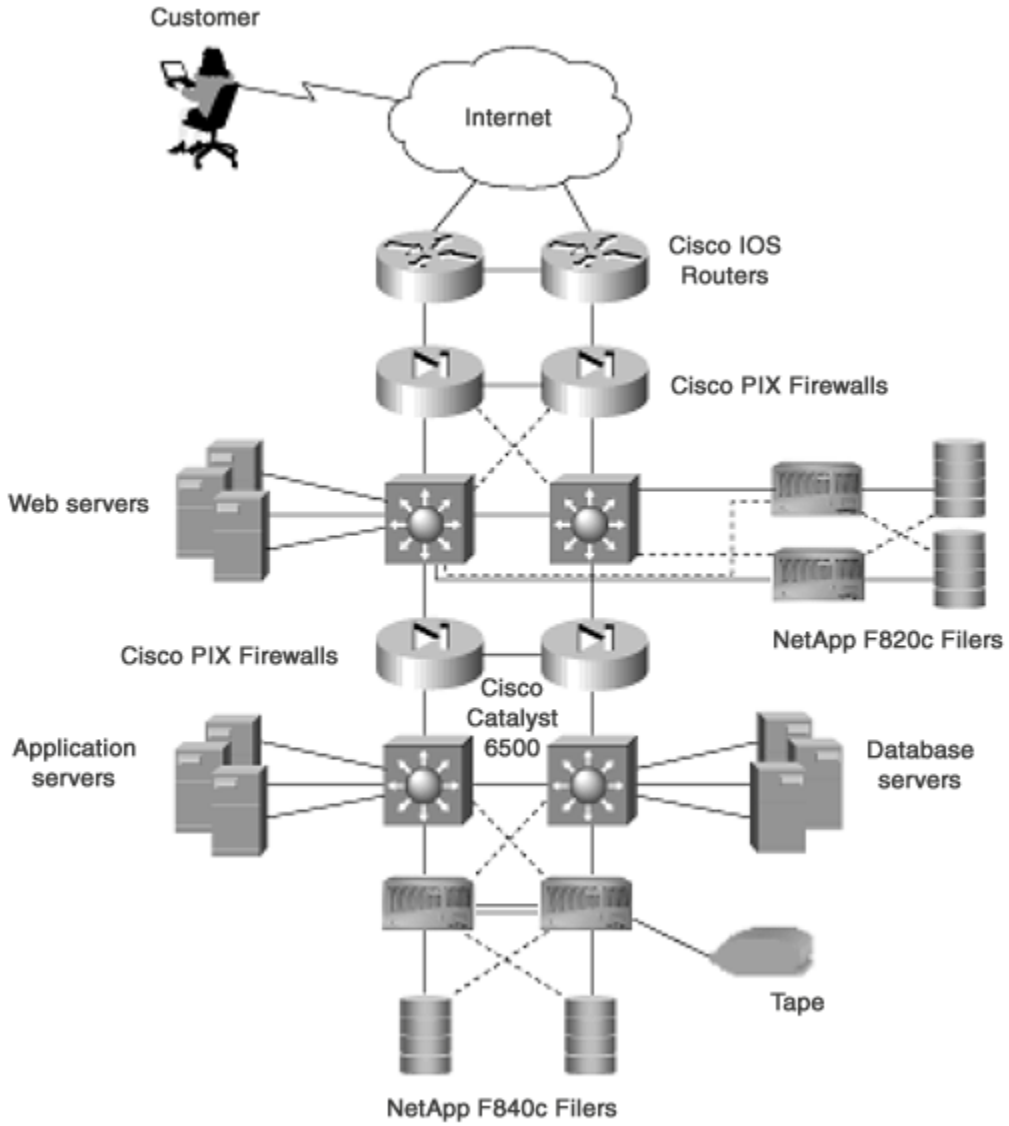

Figure 4.1. Internet E-business Application

In the direct-attached storage configuration, a redundant copy of the content is stored on each Web server, resulting in wasted storage resources. More importantly, as content changes, updating and managing content across these Web servers becomes unmanageable, leading to efficiency bottlenecks. This approach to Web content storage is inefficient in terms of scalability of the Web server farm because customers must often upgrade servers to increase storage capacity.

An alternative approach to the Web tier configuration is to consolidate multiple copies of content from Web servers onto filers. In this case the Web servers can leverage this shared resource and there is no need for data replication on the individual servers. This approach significantly reduces the content-management burden and makes the Web servers much easier to replace in case of failures or maintenance issues because all data is stored on filers.

Highly available architecture ensures 24x7 operation.
E-business applications must be available 24x7. Downtime is not an option because potential revenue, as well as market perception, is negatively impacted.

The configuration in Figure 4.1 shows a highly available, fault-tolerant architecture. At the Web tier, a clustered pair of NetApp F820 filers provides a highly resilient storage platform for the Web content accessed by Microsoft® Windows®, UNIX®, or Linux® based Web servers through CIFS or NFS. In case of filer head failure, data service is migrated rapidly and seamlessly to the second filer. The Web servers are aggregated at this tier using Cisco Catalyst 6500 switches and are connected through the Gigabit Ethernet ports on the switches. To ensure no single point of failure, the switches are configured with redundant network modules, power supplies, and so forth.

The back-end NetApp 840c filers, which store the application and database content, are also configured in a cluster. Highly resilient network connectivity for the application and database tiers is achieved using redundant Cisco Catalyst 6500 switches through Gigabit Ethernet connections. The Cisco Catalyst switches are configured with the PortFast feature on the ports that connect the filers and the UplinkFast feature to connect to the front-end switches. This provides rapid convergence of the spanning tree if one of the network links goes down. The network switches are configured to run Hot Standby Routing Protocol (HSRP) between them. If a switch breaks down, the other switch takes over the load of the failed switch in a few seconds without any observable interruptions at the application layer.

For disaster recovery, a mirroring software feature called SnapMirror® is used to replicate databases over a secure WAN connection (see section 4.4., "Distributed Enterprise Storage with Secure WAN Connectivity," for more details) to a remote location.

Note: A typical Internet/e-business architecture also involves other elements in the Web tier such as firewalls, caching solutions, and local or global load-balancing devices. However, this paper focuses on the network- attached storage applications and implementation. Details for these other technologies, although important, are beyond the scope of this paper.

Delivers high scalability to meet escalating growth demands.
Data storage requirements for Internet and e-business applications grow at a faster pace than for any other environment. As demand escalates, the network and storage infrastructure must scale rapidly with no downtime.

The architecture shown here is scalable both vertically and horizontally. The storage capacity on the NetApp F840 filers can scale vertically up to 12 terabytes (TB) by adding disks or shelves of disks as needed, without filer downtime or reconfiguration. Because the filer storage is deployed

in a flexible, Cisco Catalyst based switched fabric, total storage capacity can be scaled indefinitely by adding more filers to the fabric (horizontal scaling).

Ensures fast response time for demanding workloads.
Faster online response for users is one of the key drivers for customer retention and loyalty in an e-business environment. For example, the "eight-second-rule" states that if a Web page does not completely load within eight seconds, customers might not return to the Web site, leading to lost revenue. While performance is typically addressed at several levels in the infrastructure design, networked storage plays an important role in meeting overall performance goals.

A NetApp filer can deliver high performance over the network because it is dedicated to reading and writing data efficiently and does not have the overhead associated with a general-purpose operating system (OS). The Write Anywhere File Layout (WAFL™) file system and NetApp Data ONTAP™ software are optimized for efficient writes and fast retrieval of data, delivering superior performance. Also, the filers offload I/O activity from application servers, improving overall system performance. Because the Redundant Array of Independent Disks (RAID) and WAFL subsystems are tightly integrated, the WAFL write-anywhere design allows scheduling of multiple writes to the same RAID stripe whenever possible, for faster operation.

Provides simplified administration that delivers low TCO.
The filers in this scenario can be remotely managed and monitored through the NetApp DataFabric® Manager Web-based management interface or by using a systems management framework that supports a standard Simple Network Management Protocol (SNMP) MIB interface such as HP OpenView. DataFabric Manager offers a unified view of groups of filers that can be drilled down to configure individual filers using the FilerView® Web-based management interface. Data backup from filers to a locally attached tape drive can be remotely administered using the Network Data Management Protocol (NDMP).

This solution offers an exceptionally low TCO by significantly reducing operational costs, complexity, and deployment time. For example, Snapshot™ technology can reduce the backup window of production databases dramatically from hours to minutes. Snapshots are frozen images of production data. Snapshots of live databases can be taken in seconds and then backed up to tape. Moreover, filers eliminate the need for database administrators to perform common administrative tasks such as disk layout tuning and retuning, allowing them to focus on value-added tasks such as delivering additional features and fine-tuning applications.

Because this architecture is based on IP-centric appliances, deployment complexity is considerably reduced, enabling e-businesses to achieve faster time-to-market for their product and service offerings.

4.2. Business Applications in the Data Center

Common enterprise business applications, such as enterprise resource planning (ERP), customer relationship management (CRM), and supply chain management (SCM), that are deployed in enterprise data centers require a highly available networked storage infrastructure. As businesses implement Internet strategies for enhanced productivity, these applications evolve into Web-enabled solutions. As a result, the architecture in this scenario can be extended with other features such as firewall security, load balancing, and so forth, to be equally applicable in an Internet-centric environment. For more details on this extension, refer to section 4.1., "Internet and E-business Applications."

Solution Description
The primary focus of an ERP environment is high application data availability. The content associated with such enterprise business applications represents a key enterprise asset, and

downtime for mission-critical applications results in significant negative business impact. The desired high availability is achieved by building redundancy into the server, network, and storage designs. Figure 4.2. shows a data center that deploys clustered NetApp F840c filers and various servers in a Gigabit Ethernet network infrastructure to support an ERP application.

Delivers a highly resilient architecture for application and database tiers.
The ERP application in this example is hosted on clustered application servers. The application is layered over a database, which is hosted on a second (database) server cluster. A NetApp F840c clustered filer configuration provides networked storage for the ERP application database. The filers, as well as each of the servers, are connected to two Cisco Catalyst 6500 switches for network link redundancy. This server access layer connects to the core/distribution layer that provides connectivity to the enterprise network infrastructure.

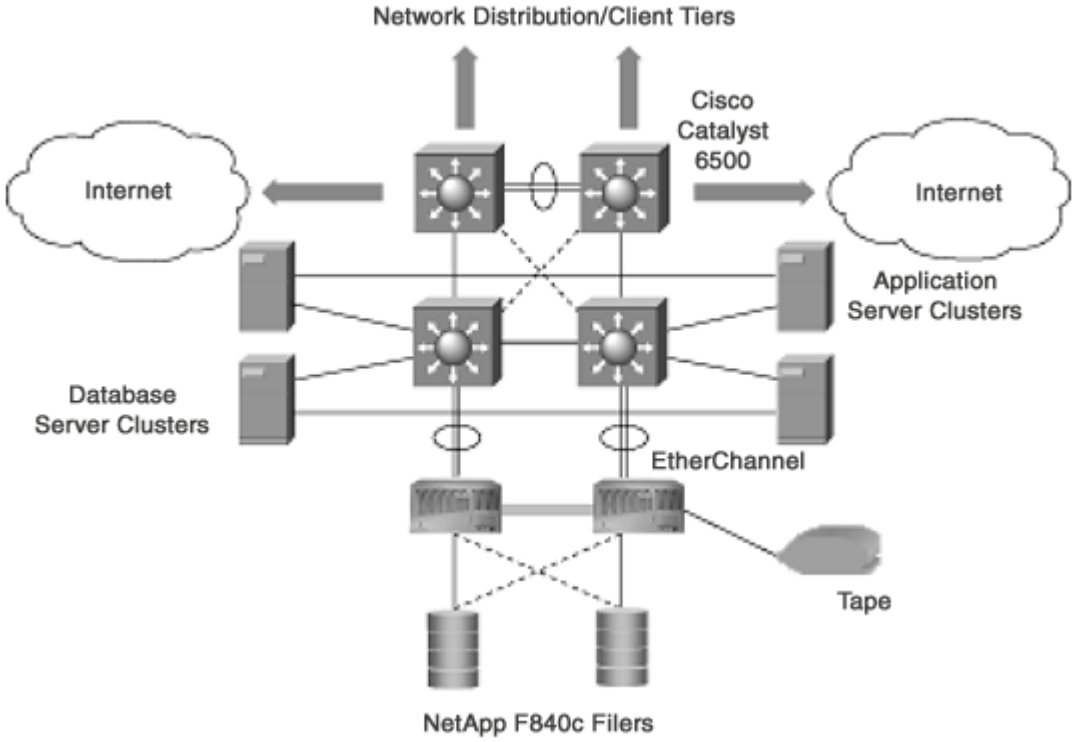

Figure 4.2. Business Applications in the Data Center

To support application availability, each application server cluster node is configured to take over the ERP application from its peer in case of server or network failure. Similarly, the database server cluster is configured so that any server can take over database services from its peers in case of server failures.

Each server cluster node is connected to two Cisco Catalyst 6500 switches through Gigabit Ethernet links to provide a redundant network path in case of link failure. HSRP is configured on the two Cisco Catalyst 6500 switches to provide automatic failover if one of the switches fails. The NetApp F840 clustered pair of filers can take over each other's data services in case of the individual filer head failure. This provides full, active-active, application-level redundancy to the overall architecture.

Additional redundancy is built into storage networking by using dual-homing for the NetApp F840 filer cluster nodes. Dual-homing (also known as virtual interfaces, or VIFs) provides additional availability and scalability to the storage network design.

In standard dual-homing mode (single-mode VIF), the interfaces are in active/standby configuration. If one of the links goes down, it triggers the standby interface to take over servicing network requests. The filer's Data ONTAP appliance OS monitors the active link and handles the failover.

Multimode VIFs can be configured to use EtherChannel® capabilities that are supported by Cisco Catalyst 6500 switches, which allows all interfaces to be active. This capability enables them to act as a single virtual interface (with combined bandwidth of multiple individual links now available). The switch monitors the status of the links and handles link failures.

The PortFast feature is enabled on the Cisco Catalyst switches on the ports that connect the filers, which allows faster convergence of the spanning tree and switch-over to alternative links for network traffic. The Cisco Catalyst 6500 switches can also be configured to join ports on two different modules into an EtherChannel configuration. This configuration eliminates downtime if one of the links or modules fails.

For efficient use of network bandwidth, servers are connected to separate virtual LANs (VLANs) that are configured on the Cisco Catalyst switch that limits traffic on the Ethernet subnets. The filers are in their own separate VLAN on the back end, and the Cisco Catalyst switch routes traffic intended for them to that subnet. Gigabit EtherChannel trunking provides higher aggregate bandwidth to networked storage required by some ERP databases.

Delivers exceptional data recoverability and simplified administration.
Data represented in a business application must be easily and quickly recoverable. According to a recent Gartner study[2] roughly 40% of the unplanned downtime is caused by corrupted data due to application errors. To reduce administration overhead, the ERP applications and databases can be structured so that development, test, and production environments are hosted on separate volumes. This approach allows features such as VolCopy to move program and data components around easily.

For a stable production environment, testing is often required before database applications are introduced to the production environment. Snapshots can be used to test changes to the production database. In this scenario, Snapshots of database table spaces can be taken and tests can be run. If the test fails, the database can be restored to a "known good" state using the SnapRestore™ facility within minutes, without reverting to tape archives.

Periodically scheduled Snapshots can help recover from corruption due to application errors in a production database. Multiple Snapshots of ERP data are available online and the entire ERP data can revert in minutes to a previous version using SnapRestore. Snapshots dramatically reduce the backup window of databases from hours to minutes and increase application availability. Snapshots of terabyte-sized volumes containing ERP database table spaces can be created in seconds and backed up to tape without interrupting the production database. Databases can be remotely replicated over a WAN for disaster recovery (see section 4.4., "Distributed Enterprise Storage with Secure WAN Connectivity," for more information). These features dramatically improve application availability and data recoverability in the solution architecture.

Delivers vertical and horizontal scalability.
Another key benefit offered by this solution is that storage can be scaled independently of the database or application servers without any downtime to the application. Shelves of disk space

can be added as needed to an F840 cluster to scale vertically to 12TB. Therefore, storage capacity can be effortlessly added to the volume supporting an ERP database while the database and filer are in operation. To accommodate growth, networked storage can be scaled horizontally by deploying additional NetApp F840s. Volumes can be moved between filers for load balancing using VolCopy.

Lowers TCO.
Finally, this architecture, which is based on networked storage for ERP databases, reduces operational cost, complexity, and deployment processes dramatically. This approach significantly reduces acquisition and maintenance costs, and therefore the TCO of the ERP application.

4.3. Workgroup Collaboration

Engineering environments such as software development or engineering design groups demand an infrastructure that enables data sharing over a high-performance LAN with low management overhead. Applications with these requirements include EDA simulators and synthesis tools, and tools for CAD, CAM, CAE, source code management, version control, and so forth.

In addition, there is a critical need within an enterprise to consolidate multiple files from Windows NT® user desktops onto a data-consolidation platform such as network-attached storage appliances to enable data sharing among users (termed NT consolidation).

Solution Description
The architecture in Figure 4.3. effectively addresses the requirements in the workgroup collaboration environment.

Highly resilient architecture provides high availability.
Various types of clients are aggregated at the client access tier using redundantly configured Cisco Catalyst 4000 switches through 10/100 Fast Ethernet ports. Redundant Gigabit Ethernet uplinks connecting to the distribution layer switches ensure no single point of failure by providing alternative paths, fast convergence in case of link failures, and load balancing across these redundant switches.

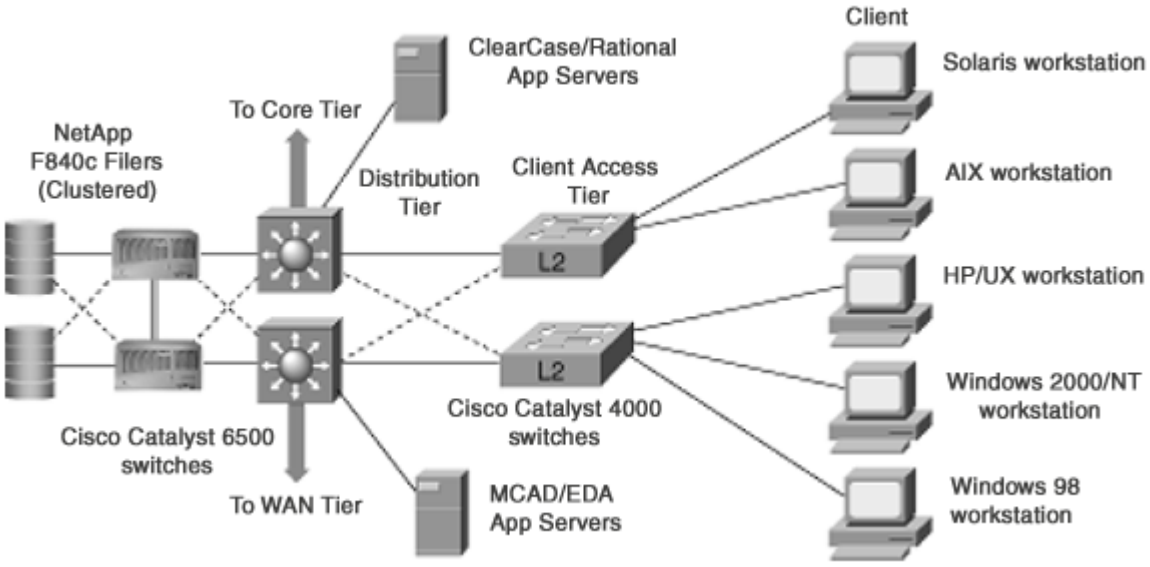

Network Appliance Inc.

Figure 4.3. Workgroup Collaboration

A clustered pair of NetApp F840c filers ensures that data service is automatically migrated over to the second head, in the unlikely case of a filer head failure. The filers are connected through Gigabit Ethernet ports on the Catalyst 6500 multilayer switches at the distribution tier, which also provides connectivity to the application servers in case the applications used by the workgroup requires them at this tier. This creates a private dedicated storage network for connecting the application servers that run engineering applications to the filers.

To extend high availability, these distribution layer switches are also deployed in a redundant configuration. They are configured with HSRP, which allows one switch to assume the load of a failed switch within seconds without any interruption to data connectivity at the higher layer. The application servers-in the case of the engineering environment-are also connected to these switches to allow data sharing with other organizations.

Delivers seamless data recoverability.
Software code and engineering designs are the core assets of any engineering organization. Data corruption leads to rework, which negatively affects time-to-market. The solution architecture described in this section allows seamless data recoverability. The Snapshot, SnapRestore, and SnapMirror features of the NetApp Data ONTAP software ensure that data in the production volume is available without interruption. Periodic Snapshots can be automatically scheduled at the desired set frequency. Snapshots of terabyte volumes can be taken in seconds and provide up to 255 read-only images of data online for immediate recovery.

Snapshots help individual users restore personal data files in case these users need to go back to a previous version, without resorting to help from their IT organization. If the application data is corrupted, SnapRestore allows restoration of complete volumes to a specific Snapshot, which avoids hours or days of data restoration from tape.

Engineering environments tend to be complex to administer. Simplification of data management is key to smooth operation. Core administrative requirements include:

- Ability to back up without impacting the production system
- Simple processes for system upgrades
- Ability to recover data effortlessly during a data corruption or data loss event
- Simplified disaster recovery implementation

For backup and quick disaster recovery, data can be remotely replicated over a WAN using the SnapMirror feature. The data is replicated in asynchronous mode with no performance penalty or WAN link hogs (refer to section 4.4., "Distributed Enterprise Storage with Secure WAN Connectivity," for more details). Data from filers can be backed up to a central tape library that supports the NDMP protocol over the IP network. Backups are performed from Snapshots so that the production volume can remain online, dramatically reducing the backup window. Thus, the NetApp filer's Snap product family ensures high data availability and recoverability of key files.

Serves heterogeneous operating system environment.
NetApp filers enable multiple operating system platforms to share the storage resources. Data ONTAP, the microkernel within NetApp filers, includes native support for industry-standard protocols such as NFS and CIFS. Such multiprotocol OS support allows true data sharing, where a single copy of data can be shared between Windows and UNIX clients. Therefore, the clients in this environment can be a combination of UNIX or Windows NT machines running engineering

applications that are connected through the network access layer. The SecureShare® feature within the microkernel provides file locking for safe sharing of data in such heterogeneous environments.

Delivers industry-leading performance.
NetApp filers deliver high performance over the network because they leverage an appliance architecture that is dedicated to reading and writing data efficiently and does not have the overhead associated with a general-purpose OS. WAFL file system and Data ONTAP are optimized for efficient writes and fast retrieval of data, ensuring minimum head seeks. Also, the filers offload I/O activity from application servers, improving overall system performance. Because the RAID and WAFL subsystems are tightly integrated, WAFL knows the current load and usage of every disk in the array. This allows WAFL to automatically tune its performance for a given workload.

On the network layer, the switched VLAN topology allows full-duplex 10/100 Ethernet bandwidth between the clients and the switch as well as a dedicated Gigabit Ethernet connection to the NetApp filers for storage access through the distribution tier.

Delivers exceptional horizontal and vertical scalability.
The solution described above provides scalability options including adding more clients at the network access layer through additional switches connecting to the distribution layer. The storage capacity on the NetApp F840c filers can scale vertically up to 12TB as needed with no downtime or reconfiguration. The architecture can also scale horizontally by adding filers to the network in a few minutes.

Provides simple-to-use management.
The overall distributed storage environment can be managed centrally using DataFabric Manager (Web-based GUI). The network components and filers can also be integrated into major systems management platforms with built-in SNMP support.

4.4. Distributed Enterprise Storage with Secure WAN Connectivity

Organizations that have geographically distributed sites, such as offshore development offices that are connected to a central headquarters or to regional offices, demand distributed storage networks for effective collaboration among geographically dispersed groups as well as disaster recovery implementation. These remote sites are typically connected to each other and the central site through a hub and spoke topology and therefore require a secure WAN environment for data exchange across the links. The storage network must be designed for safe and efficient WAN bandwidth use.

Solution Description
Figure 4.4 shows an example of how an enterprise can deploy distributed storage environments to address the requirements described above. This scenario extends the previous engineering workgroup scenario to include an offshore development group that needs to work collaboratively with the core engineering team at the headquarters site. The typical applications include sharing files with the local team as well as having read-only access to the central engineering folder.

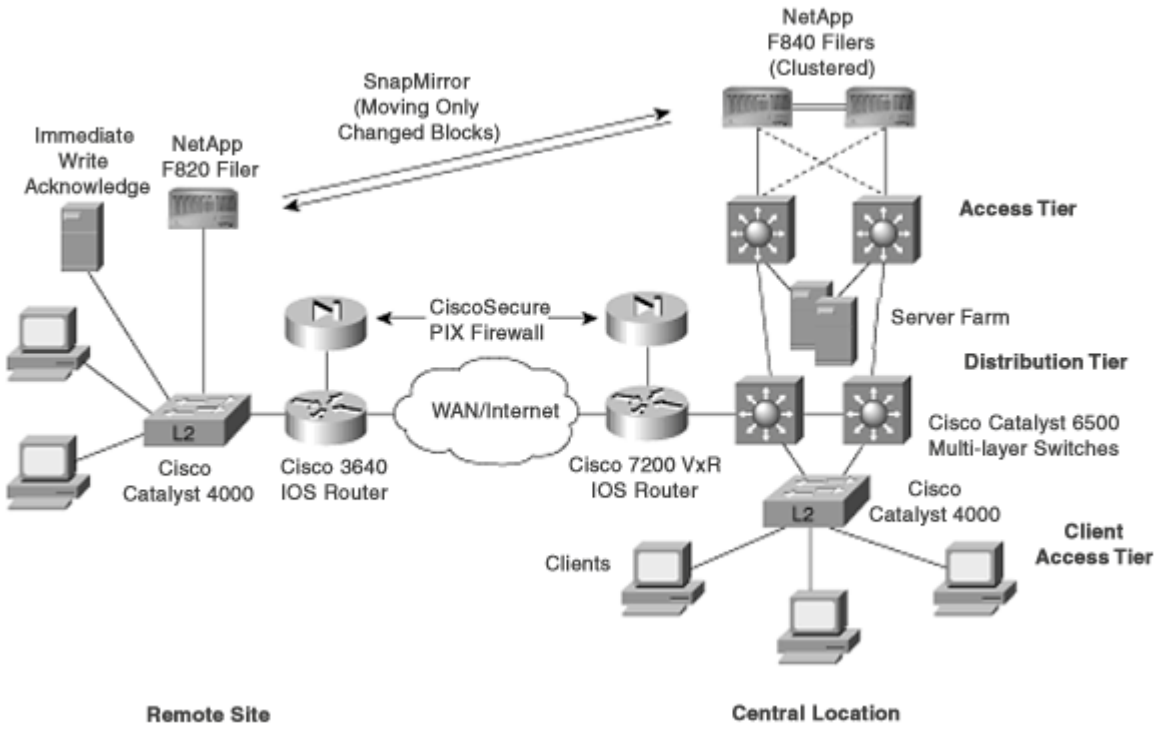

Figure 4.4. Distributed Storage over Secure WAN

Delivers proven security and enhanced quality of service.
The WAN link between an engineering site at a central location and remote offices must be secure before any sensitive data files are exchanged or copied across it. The WAN link must also be efficiently utilized for high return on investment (ROI) from the cost of the link. The WAN connectivity in this scenario can be achieved through a dedicated link (using WAN services such as Frame Relay, ATM, and so forth) connecting to the WAN module of the central site infrastructure. An alternative approach is to deploy secure VPN connectivity over the data links that are already in place.

Cisco IOS® routers such as the Cisco 7200 and Cisco 3640 routers in the figure provide various WAN modules, Cisco IOS firewall security, and hardware-accelerated encryption for virtual private network (VPN) connectivity. The Cisco IOS software also comes with various QoS features such as Committed Access Rate (CAR) and Weighted Fair Queuing (WFQ), which address traffic congestion issues at the WAN link. These features ensure that the WAN link bandwidth is appropriately used by storage-related traffic and does not drain resources from other mission-critical application data.

Provides efficient data replication.
Collaboration among geographically dispersed groups demands data sharing among different locations. The engineering group at a central location in this example might need to locally access an offshore group's design data. This must be achieved without compromising the data-sharing requirement within the offshore workgroup environment.

For disaster recovery, NetApp's Data ONTAP leverages its Snapshot capability to provide an automated file system mirroring capability called SnapMirror. Using SnapMirror technology, filers can replicate one or more data volumes to another filer. SnapMirror transmits only changed

blocks to the destination, utilizing minimal bandwidth. The timing and frequency of the updates can be easily controlled from a single location. The key differentiator of SnapMirror is its tight integration with the NetApp WAFL file system and Data ONTAP kernel. As a result, the data at the destination is always in a consistent state. SnapMirror can also replicate data from one filer to multiple filers of any type, eliminating the need to have exactly the same filer configuration on both sides.

The network connectivity at the remote site is typically addressed using LAN switches such as the Cisco Catalyst 4000 Series. Similarly, Cisco Catalyst 6500 Series switches connect the server farm with back-end filer storage using Gigabit Ethernet connectivity at the central site. The Cisco Catalyst 6500 switches are configured with features such as VLANs and EtherChannel, and high-availability features including HSRP, fast routing protocol convergence, and redundant switch modules.

Delivers modular architecture for scalability.
More remote sites can easily be added for WAN and VPN connectivity to the Cisco 7200 routers. For increased data protection, dedicated Cisco Secure PIX® firewalls are deployed to offload the security functionality from Cisco IOS firewall-enabled Cisco 7200 routers. The modular Cisco Catalyst 6500 platform allows various modules to be added to the same chassis, enabling the user to avoid fork-lift upgrades at the network layer.

The solution architecture shown in the example provides seamless scalability. To add more storage at the central site, customers can easily add additional filers or disk storage to the existing infrastructure. The NetApp F800 series filers can support up to 12TB of storage and also support clustered configurations. This enables the solution architecture to provide much higher availability of storage resources.

Simplified management of distributed networked storage components.
As in other solutions described in this document, DataFabric Manager (Web-based GUI) provides Web-based management for all filers in this architecture. Similarly, SNMP support enables the network components and filers to be integrated into major management platforms.

5. Conclusion

Enterprises face considerable challenges brought by the rapid adoption of emerging Internet business applications and the associated storage infrastructure requirements. Storage networking aims to deliver solutions to significantly mitigate these challenges. Network Appliance and Cisco Systems, the leaders in storage and networking solutions, are collaborating to deliver proven networked storage solutions for the enterprises to meet these challenges. These solutions deliver all the key business benefits such as scalability, performance, simplified management, availability, and security, while leveraging existing investments and expertise in IP networks and drastically reducing TCO.

6. Appendix A: Cisco Systems Products

Catalyst 6000 Series Switches
The Catalyst 6000 Series delivers exceptional scalability and price/performance multilayer switching solutions for campus networks. It is designed to address the increased requirements for gigabit scalability, high availability, and multilayer switching in backbone/distribution and server aggregation environments. The Catalyst 6000 Series provides intelligent server-switching and Web-scaling technology, integrated voice solutions, and LAN/WAN integration.
http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/

Catalyst 4000 Series Switches
The Catalyst 4000 Series provides an advanced enterprise switching solution for wiring closets and data centers. The Catalyst 4000 Series integrates the features and performance required in today's wiring closets and provides the flexibility for future enhancements.
http://www.cisco.com/warp/public/cc/pd/si/casi/ca4000/

Cisco 7200 Series Routers
The Cisco 7200 Series router is an enterprise WAN gateway platform for power branch connectivity. It provides high performance, density, modularity, and availability with low per-port prices, offering fully integrated network services. The Cisco 7200 Series router provides a scalable multiservice solution for enterprise central sites to integrated data, voice, and video networking. The Cisco 7200 Series is an ideal site-to-site VPN router choice for customers who seek high-performance, hardware-assisted tunneling and encryption services that are suitable for private WAN and VPN applications.
http://www.cisco.com/warp/public/cc/pd/rt/7200/

Cisco 3600 Series Routers
The Cisco 3600 Series offers high-performance, modular, multiservice access platforms for medium and large offices and smaller Internet service providers (ISPs). It protects customers' investments in network technology and integrates the functions of several devices into a single manageable solution. The Cisco 3600 Series provides solutions for data, voice, video, hybrid dial access, VPNs, and multiprotocol data routing.
http://www.cisco.com/warp/public/cc/pd/rt/3600/

Cisco Secure PIX Firewalls
The Cisco Secure PIX Firewall provides full firewall protection that conceals the architecture of an internal network from the outside world. All Cisco Secure PIX Firewall models have built-in IP security (IPsec) encryption, permitting both site-to-site and remote access VPN deployments. These firewalls operate on a hardened operating system that protects the security of the device and the network. Cisco Secure PIX firewalls offer stateful session inspection, user authentication and authorization, and technology to allow for stateful failover should a path in the network fail.
http://www.cisco.com/warp/public/cc/pd/fw/sqfw500/

Other Cisco Network Product Offerings

http://www.cisco.com/public/products_prod.shtml

Cisco Definitions
For definitions of terms used in this document, refer to the Cisco online Dictionary of Terms and Acronyms at:
http://www.cisco.com/univercd/cc/td/doc/product/cable/cab_rout/cmtsfg/ufg_pref.htm#xtocid5

7. Appendix B: Network Appliance Products

NetApp F840c and F840 - Large-Scale, Clustered Enterprise Filers
For maximized data availability in mission-critical environments, the NetApp F840c clustered filers and F840 filers meet the open storage challenges for today's large enterprises. NetApp F840c filers provide industry-leading performance and capacity, scaling up to 12TB of network-accessible data. The F840c architecture integrates multiple F840 filers in an active-active clustered failover configuration to provide high availability and scalable performance to multiple networks. With throughput and response times that support the most demanding applications, the NetApp F840 filers provide an ideal solution for enterprise consolidation. The NetApp F840 filers deliver high performance, high availability, and effortless management. Scaling to 6TB (12TB in

clustered configurations), the NetApp F840 brings enterprise data consolidation to a new level.
http://www.netapp.com/products/filer/f800_ds.html

NetApp F740 - Entry-Level Enterprise Filers
The NetApp F740 ideal entry-level enterprise filer delivers the power to boost productivity in departmental applications, making it the optimal solution for software development, CAD, simulation applications, and medium-sized ISPs. The NetApp F740 scales to 1TB, providing consolidation benefits to any medium-sized business with productivity and access requirements for hundreds of users.
http://www.netapp.com/products/filer/f740_ds.html

NetApp F85 - Remote Office/Workgroup Filer
The NetApp F85 is an economical workgroup file server with enterprise-class features, availability, and multiplatform capability at an economical price. Architecturally similar to powerful NetApp enterprise-class filers, the NetApp F85 is available in free-standing or rack-mount configurations. The NetApp F85 scales to up to 648GB and is ideal for remote offices and branch offices as well as small- to medium-sized workgroups and ISPs.

Data ONTAP and WAFL
Data ONTAP is the real-time microkernel-based OS upon which all NetApp appliances are based. NetApp's architecture for data availability and management is based on the foundation of Data ONTAP. WAFL is the Data ONTAP file system that decides how and where files are written to disk blocks. The highly advanced, patented, file-system-based RAID architecture of WAFL provides disk failure resiliency without the performance penalties normally associated with general-purpose parity-based RAID designs.
http://www.netapp.com/products/software/ontap.html

Snapshots
Snapshots are read-only images of a volume that Data ONTAP can automatically create at different points in time. Snapshots are efficiently created from data pointers instead of copied blocks because of unique Data ONTAP copy-on-write policy. Snapshots are accessible to individual users and allow them to independently restore files. Snapshots are a standard feature of Data ONTAP.

SnapRestore
SnapRestore is a standard feature of Data ONTAP that allows a volume to be restored to a Snapshot corresponding to its state at a previous point in time. Terabytes of data in a volume can be restored to a "known good" Snapshot in minutes.
http://www.netapp.com/products/software/snaprestore.html

SnapMirror
SnapMirror is an add-on feature of Data ONTAP that allows remote replication of volumes over a WAN. A volume on a filer can be safely and efficiently replicated to a read-only volume on a remote filer over a WAN using SnapMirror. SnapMirror uses Snapshots to track changed blocks in a volume and transmits them asynchronously to the remote filer.
http://www.netapp.com/products/software/snapmirror.html

Clustered Failover (CFO)
CFO allows two filers to be connected in a highly available configuration where data services owned by one filer head can be taken over by its partner in an active-active setup. CFO requires a high-bandwidth, low-latency interconnect using ServerNet and shared connections to disk shelves between filer heads. CFO delivers 99.99 percent availability.
http://www.netapp.com/products/software/clustered.html

DataFabric Manager (DFM)
DataFabric Manager is a new unbundled software application for managing NetApp appliances, including filers and filer clusters. This product allows storage and network administrators to centrally monitor and remotely configure groups of appliances and volumes. DataFabric Manager provides a unified administrative view of the total storage network of NetApp appliances.
http://www.netapp.com/products/software/datafabric.html

FilerView
FilerView is a Web-based interface for remote filer administration. FilerView enables remote management and monitoring of filers from a browser on a remote machine. FilerView is a standard feature of Data ONTAP.

[1] Source: IDC, 2000 Disk Storage System Forecast & Analysis (1999-2004), 12/2000
[2] Gartner IT Symposium 2000

NetApp®

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089
www.netapp.com