NetApp MultiStore™ and SnapMover®

Adding Business Value with Storage Consolidation, Data Availability, Security, and Load Balancing

Miroslav Klivansky and John Phillips | Network Appliance | January 2005 | TR 3150

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.

Abstract

NetApp MultiStore allows a filer's storage and networking resources to be partitioned into multiple virtual filers, each of which appears as an individual filer on the network. The dynamic, flexible nature of virtual filers simplifies UNIX and Windows storage consolidation, and offers unified access to data over both CIFS and NFS. Virtual filers can be wholly mirrored to other filers over a network for data migration, disaster recovery, or load balancing purposes. Virtual filers can be grouped into private network address spaces and belong to different security domains for maximum security. MultiStore simplifies data storage and management for companies or service providers requiring highly available storage and secure, multi-domain flexibility.

Table of Contents

Table of Figures

1) Introduction

Reacting to rapidly changing business needs while finding efficient ways to store and manage data has changed the storage landscape. Data storage systems must be able to scale. They are expected to provide reliable and highly available service. They must be flexible to continually adjust to changing business needs. They must be simple to deploy and manage, reducing administrative costs and delivering substantial returns on investment.

NetApp® filers are storage appliances running Data ONTAP™ software. Filers offer industry-leading reliability and performance, and simplify many otherwise difficult data storage and management tasks. File systems can be scaled quickly and easily with no downtime, allowing customers to adapt and grow in ways that best suit their evolving needs. Filers run an optimized microkernel, and provide transparent, shared access to data using industry-standard protocols for Windows®, UNIX® and Web clients.

Using NetApp MultiStore technology, a filer's networking and storage resources can be partitioned into virtual filers, referred to as a Vfiler™. Each Vfiler appears on the network as a separate filer. More importantly, virtual filers and the resources they use are dynamically configurable, can be replicated to other filers using NetApp SnapMirror® technology for disaster recovery purposes and data migration, and in special configurations instantly migrated to other filers without copying data by using SnapMover® technology. Virtual filers can be grouped into one or more IPspaces, each of which represents an independent, secure network with its own addressing and routing. These features allow companies to architect and manage their IT infrastructure more effectively. They also give service providers the tools to quickly and efficiently provision new customers, and organizations the confidence to outsource complex IT management tasks to a single, cost-effective center of excellence.

This paper will describe how NetApp MultiStore technology enables companies to better manage, consolidate, migrate, and replicate critical data with minimal effort and maximum return.

2) MultiStore Solutions

2.1) MultiStore Overview

The filer hardware is made up of CPU(s), network cards, Fibre Channel controllers, power supplies, disk drives, etc. Using MultiStore, a filer's storage and networking resources can be effectively partitioned and dynamically assigned to virtual filers. This virtualizes the physical resources and moves beyond the logical architectural limitations inherent in a single physical filer. Up to 32 virtual filers (in addition to the hosting filer, known as vfiler0) can be created and hosted on a NetApp filer, each serving data as a filer does. Other network computers communicate with virtual filers as they would with any other file server.

Vfiler facts:

- Virtual filers appear on the network as discrete multiprotocol file servers, and are accessed as any other Windows or UNIX file server would be
- Companies can deploy a Vfiler into full production in a matter of minutes with only a few software commands
- Each Vfiler independantly controls access to its storage resources according to its security settings and permissions assigned to groups and users by Windows and/or UNIX system administrators
- Network settings such as IP addresses, IPspace associations, DNS information, Windows domain and other settings are explicitly assigned to each Vfiler
- Virtual filers and their resources can be replicated or moved to other filers, dramatically simplifying data consolidation, migration, and disaster recovery

Resources can be added, removed, or moved between virtual filers at any time. Figure 1 below depicts a filer that also hosts three virtual filers. The virtual filers are assigned resources from the hosting filer, which still functions as a normal filer. Thus, the filer below appears on the network as filer1, filer2, filer3, and filer4.

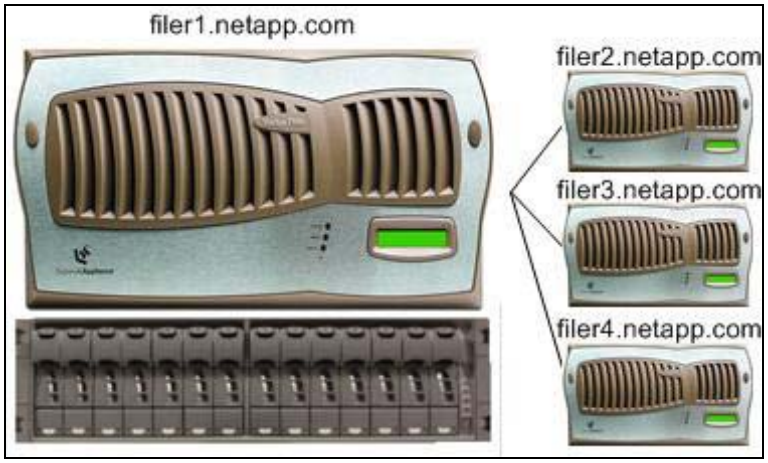

Figure 1: Hosting filer and three virtual filers

In the simplified example above, the hosting filer and each Vfiler is a member of the same UNIX NIS domain "netapp.com", the same Microsoft® Active Directory domain netapp.com, and the same DNS domain netapp.com. However, each Vfiler could be created in a separate IP space and assigned membership in different security domains and DNS namespaces. More information about Vfiler networking features will be discussed later in this document.

Business Value of Virtual Filers

Different administrators, groups, or organizations may be responsible for managing the hosting filer and/or each Vfiler. The hosting filer and certain aspects of each Vfiler can all be managed independently. This allows one department or organization to manage the hosting filer, while allowing others to independently and securely manage their own data.

The hosting filer administrator or storage engineer is responsible for managing volumes, disks, RAID groups, Snapshots™, and Backup. A Vfiler administrator is responsible for managing Windows and/or UNIX security, file system permissions, and antivirus settings for their organization's Vfiler. Vfiler administrators for different organizations do not have permission to access other virtual filers unless explicitly allowed. This allows for more flexible management policies and greater security.

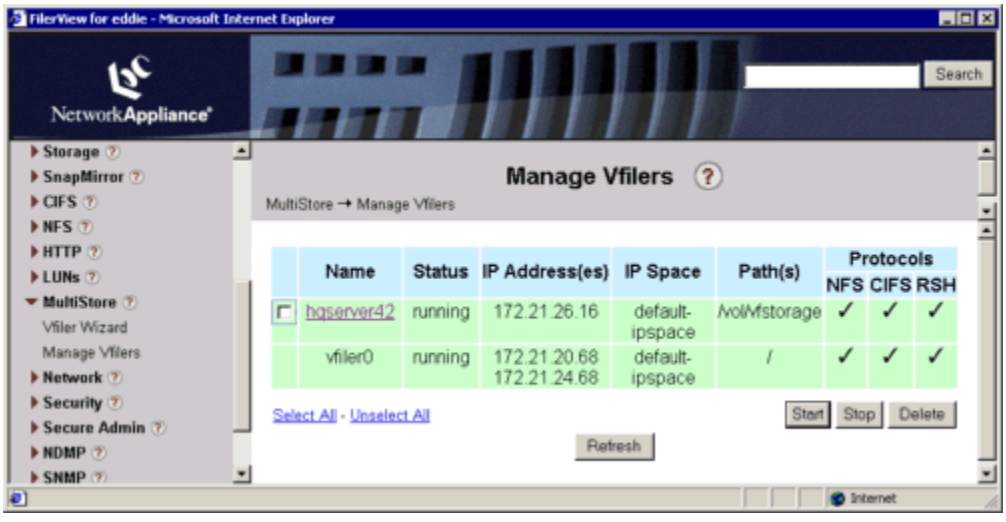

Figure 2: Managing MultiStore with FilerView

Virtual filers may be managed using the Command Line Interface (CLI) , the web-based NetApp FilerView application, or the centralized DataFabric® Manager framework. Some advanced tasks, such as those related to the SnapMirror integrated disaster-recovery and migration functionality must be managed using the CLI. Other tasks may be performed via OS-level tools like Active Directory. The Manage ONTAP™ solution also allows programmatic management via Java, C, and Perl toolkits. With management options available to meet a range of needs, it is possible to integrate filers and MultiStore into a wide range of existing operations and processes.

Virtual filers can be easily created and destroyed as needed to accommodate temporary projects or changing project priorities. It takes only minutes to deploy a new Vfiler, and even less time to destroy one that is no longer needed. No additional hardware needs to be deployed or managed. The unprecedented ease and speed opens doors to more efficient operations and project opportunities.

As business needs and workloads change, virtual filers may be migrated to ensure that critical projects receive the required level of service and response times, while retired projects are still available online via lower-tier storage. Using different filer models in the same infrastructure allows critical projects to run on the most powerful hardware, while other projects are still available via smaller models. As priorities change, the administrator can adjust the Vfiler deployments with just a few commands. This flexibility leads to greater return on investment (ROI) and lower total cost of ownership (TCO).

Storage management is virtualized, allowing the IT architect to separate many of the logical needs of the infrastructure from the physical hardware used to build it. Because less hardware is required management and acquisition costs can be reduced, and adaptability improved.

In conjunction with SnapMover technology, a Vfiler can be migrated to another NetApp filer in seconds, with no data being copied and no interruptions to the user. This is especially useful in dynamic computational environments like Grid Computing where huge computational demands change rapidly as jobs and projects start and finish. Combined with Grid technology, the load balancing capabilities provide opportunities for unprecedented application scaling efficiencies.

MultiStore and SnapMover technologies can also be used with Vfiler migration for system maintenance with no operational interruptions. Virtual filers are migrated from a filer needing maintenance to another filer. The maintenance is completed and the filer rebooted, and the Vfiler is migrated back to the original filer – all with no interruptions or need to re-establish connections with the filer if using NFS or iSCSI. With SnapMover, the migrations take just seconds and don’t require additional network resources. Non-disruptive maintenance delivers higher levels of service and allows more profitable operations by being able to offer more stringent Service Level Agreements.

SnapMirror can be used to replicate virtual filers to one or more target filers, where the mirrored virtual filers can be quickly activated for disaster recovery purposes. This significantly simplifies disaster preparedness, and reduces the chances of anything else going wrong in the stressful minutes following a disaster.

Vfiler Setup

Once a filer is installed into equipment racks, and the power, disk, and network cables are installed, configuration is exceptionally quick and straightforward. It's not unusual for experienced administrators to unpack a new filer and deploy it into production in about an hour. Virtual filers are even easier as they can be created and deployed in minutes with a few simple software commands.

Virtual filers are associated with three different flavors of data. Understanding the differences is important for understanding the rest of this paper. Let’s call the first kind of Vfiler information “metadata”. This is the information that the MultiStore software associates with each Vfiler and uses to keep track of details like the Vfiler state, what protocols are licensed on the Vfiler, what volumes are owned by the Vfiler, etc. The metadata is book-keeping information needed to have the software function correctly. The second kind of information is configuration data. This is kept in the /etc directory of the Vfiler, just like configuration data for a physical filer. The configuration data is what gets defined during the Vfiler setup process, and can be updated along the way to reconfigure the Vfiler. Configuration data typically deals with areas like network addressing, security authentication, etc. The last kind of data is the actual data that the Vfiler is storing and serving on behalf of the hosts and applications using the filer.

This is stuff like user files, database tables, etc. The stored data is kept in the Data ONTAP volumes owned by the Vfiler. Keeping these distinctions clear will help to better understand Virtual Filers.

Depending on which protocols are licensed on the filer, there are a few simple steps that must be taken to configure a Vfiler after it is created. Each Vfiler is set up like any filer. The configuration files are all created in the /etc directory of the Vfiler's primary storage unit when all of the prompts have been answered and setup is complete.

There are three main processes that apply to setting up a Vfiler:

Setup Process	Information That Must Be Supplied by the Administrator
The initial/basic setup routine	▪ Administration host(s) ▪ DNS domain name and nameserver information ▪ NIS information (if applicable) ▪ Root/administrative password
CIFS setup (for Windows environments)	▪ WINS server addresses (if applicable) ▪ Administrative/root password ▪ Type of user authentication (Windows NT4, Windows 2000, Windows .NET domain, Workgroup Authentication, /etc/passwd, or NIS-based authentication ▪ If Windows domain authentication is selected, the password with the permissions to create a domain computer account
Enabling NFS	▪ By default, the NFS protocol is not turned on or active. A single command activates NFS (NFS must already be licensed on the hosting filer)

For more information on how NetApp filers function within Windows domains, please see technical reports TR-3113, TR-3124, and other relevant documents available in the NetApp Library at http://www.netapp.com/library/.

MultiStore and Flexible Volumes

NetApp introduced Flexible Volumes and Aggregates in Data ONTAP 7G. Data ONTAP 7G delivers exceptional business value to help today's enterprise liberate its data for more robust, high-performance applications in workgroups, departments, and data centers. New, built-in software features allow storage to be pooled (and managed as a pool) in a data-centric (not disk-centric) approach without regard for physical disk locations or configurations. Like the name implies, a Flexible Volume (FlexVol) may be flexibly made larger or smaller independent of the underlying disk size. Additional FlexVol features such as instantaneous cloning make Flexible Volumes a quantum leap in storage manageability.

Flexible Volumes are allocated out of a new constructed called an Aggregate. An Aggregate is a collection of reliable RAID groups, which in turn consist of individual disks. Storage volumes constructed in the manner similar to volumes prior to Data ONTAP 7G are still available and are called Traditional Volumes.

MultiStore works with both Flexible Volumes and Traditional Volumes. All security, storage consolidation, and SnapMirror-based migration and disaster recovery functions are identical. The only exception is instantaneous SnapMover-based migrations. SnapMover uses ownership of the underlying physical disks for the migration, and since now a single disk may hold data for multiple Flexible Volumes it cannot be used to migrate a FlexVol. SnapMover can still be used to migrate Traditional Volumes. Future versions of Data ONTAP will include the ability to migrate Aggregates using SnapMover technology.

Clustering Considerations

Filers that are part of a NetApp cluster function independently during normal operation. If one filer undergoes a system failure or is shut down, the partner filer will continue to function as itself, while also accessing the failed filer's disks and assuming its identity. Each member of a cluster must have a MultiStore license to take over its partner with a MultiStore license.

When using a NetApp cluster, up to 32 virtual filers can be created on each node of a NetApp cluster. Should an outage occur on one of the clustered systems for any reason, virtual filers will automatically restart on the takeover

filer in a matter of minutes. The virtual filers hosted by the filers of the cluster are created and configured independently. That is, each filer can host a different number of virtual filers, and the Vfiler configurations on the filers can be different from each other. In takeover mode, the functioning filer takes over all virtual filers created on the failed filer. These virtual filers include the ones created and vfiler0. Therefore, for virtual filers on the failed filer to work correctly after the takeover, each network interface used by a vFiler in a cluster must have a partner interface already defined.

When deploying NetApp clusters with virtual filers, each filer's IPspaces and partner network interfaces must be configured correctly in order for the virtual filers to restart on the functioning filer after a takeover. For more information on configuring partner interfaces for NetApp clusters, please see the *Data ONTAP System Administrator's Guide*. The SAG and other Data ONTAP documentation is available online at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

Transparent Migration

MultiStore supports transparent migration. If the networking infrastructure is up to the task, it's possible to migrate a Vfiler from one filer to another within seconds, with the higher-level protocols handling any requests outstanding during the short migration time. The level of transparency depends on the higher-level protocols. MultiStore can eliminate disruption for NFS users during the migration process without requiring any service interruptions, remounts, or system downtime. IP SAN users leveraging the iSCSI protocol can also transparently migrate between filers. The only caveat is that current versions of Data ONTAP associate CHAP authentication with the physical filer, so in environment where CHAP is used the iSCSI connections will need to be re-authenticated on the destination filer. NetApp plans to address this in future versions of Data ONTAP.

Unlike NFS, CIFS is a state-full protocol, and this essentially means that the client and server need to maintain an active TCP/IP connection during all times a file is open. Windows clients wait a finite time for a response to a CIFS request, after this timeout the redirector on the client assumes the server is down and shuts down the session. This timeout can be set on the client, and defaults to the maximum possible value of 45 seconds. If client session context is not transferred within 45 seconds, it may result in data loss for clients that have CIFS files open for writes, op-locks, etc. That means the Vfiler migration must complete within the defined timeout period. While most Vfiler migration handovers complete in less than 45 seconds, migration time depends on many factors (including the networking infrastructure) and cannot be guaranteed.

2.2) Storage Consolidation

For most companies, data storage growth combined with the rapid proliferation of UNIX and Windows servers throughout the enterprise has complicated data management. Data is spread among large numbers of disparate servers, each of which must be managed and maintained.

NetApp MultiStore functionality adds even more flexibility to the already compelling advantages offered by NetApp filers when used for data consolidation. Companies with large numbers of UNIX and Windows servers have long enjoyed the benefits of migrating data to a filer, which unifies shared access to UNIX and Windows data. Integrated cross-protocol file locking for UNIX and Windows data, ease of use and management, scalability, Network Appliance™ Snapshot technology, built-in RAID protection, and file system checksums are all features intrinsic to NetApp filers. NetApp filers are fast, simple, and reliable.

Using MultiStore, an administrator can connect to a NetApp filer anywhere in the world and create, set up, and deploy a Vfiler into production in a matter of minutes.

Figure 3 shows a NetApp filer and two general-purpose UNIX and/or Windows file servers. Imagine a scenario where the engineering data is primarily being accessed via NFS, and the corporate data via CIFS. With the march of time and new applications, it becomes advantageous to access all the data via both NFS and CIFS. Trying to patch and update the general-purpose operating systems to serve both NFS and CIFS can be complex and error-prone. In addition, the general-purpose servers can be redeployed for a range of other applications if there was a more effective way to serve the files. Licensing MultiStore on the existing filer and creating a Vfiler to replace each of the older general-purpose servers provides a cost effective solution.

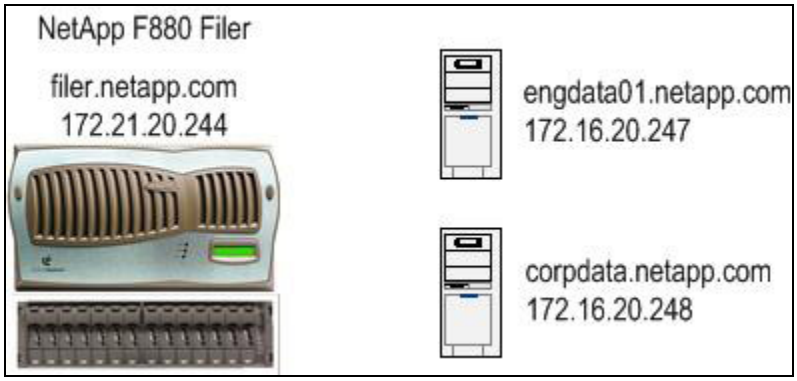

Figure 3: Data spread across NetApp filer and two general servers acting as file servers

Figure 4 shows the filer hosting two virtual filers that reuse the names of the older servers. The data previously stored on the general-purpose servers has been copied to the new virtual filers, and is now available via both CIFS and NFS. The general-purpose servers have been redeployed elsewhere, providing investment protection and greater return on investment.

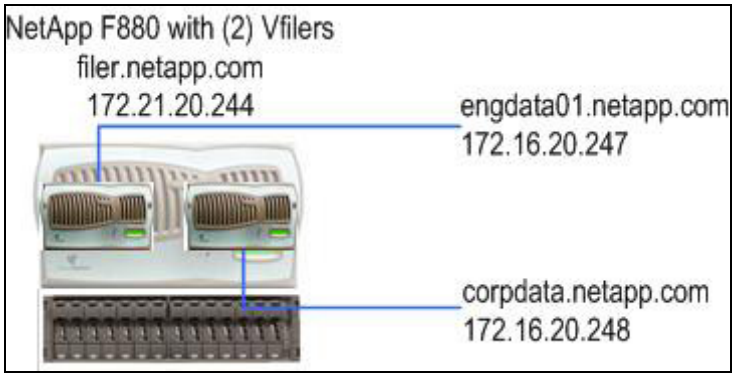

Figure 4: Data consolidated onto a single NetApp filer using MultiStore

There are many benefits to consolidating data from a large number of general-purpose servers to virtual filers running on a NetApp filer. The data can be centrally managed, backed up, easily scaled, and replicated to other filers for disaster recovery purposes using proven NetApp SnapMirror technology. In addition, powerful NetApp features like Snapshots, Flexible Volumes, and FlexVol Cloning become available to simplify management and empower the IT administrator.

In many cases, the consolidation is not readily apparent to users:

- As far as users or applications are concerned, the file servers exist as they did before, using the same names and IP addresses
- Most scripts or references to the servers or their network IP addresses within applications do not have to be changed
- Records in the corporate DNS namespace do not have to be modified by an administrator

2.3) Vfiler Migration

Through integration with either NetApp SnapMirror or SnapMover virtual filers can be migrated to other filers over a Local Area Network (LAN), Wide Area Network (WAN), or back-end storage interconnect for the purposes of load balancing or maintenance. Integrated mirror relationships may be established to simplify the migration of virtual filers to other physical filers.

Preparing for the Migrate or Disaster Recovery (DR) procedures

Note: For complete information and worksheets designed to help guide the configuration process, see the Data ONTAP 7G MultiStore Management Guide for more information.

Vfiler migration is motivated by a number of reasons. Migration may be desirable for maintenance without interruptions, workload balancing for performance, or as the only option in a disaster recovery scenario. Prior to defining a "Vfiler migrate" or "Vfiler DR" configuration it's important to verify that each of the filer(s) to be used are configured properly. Differences in the physical configuration of the destination filer, such as the number of network cards, etc., can be taken into account while establishing a migrate or DR setup. Some things however, such as the name(s) of the filer volumes must match in order for SnapMirror to function properly.

Figure 5 below summarizes the steps which should be completed prior to defining a SnapMirror migration or disaster recovery relationship.

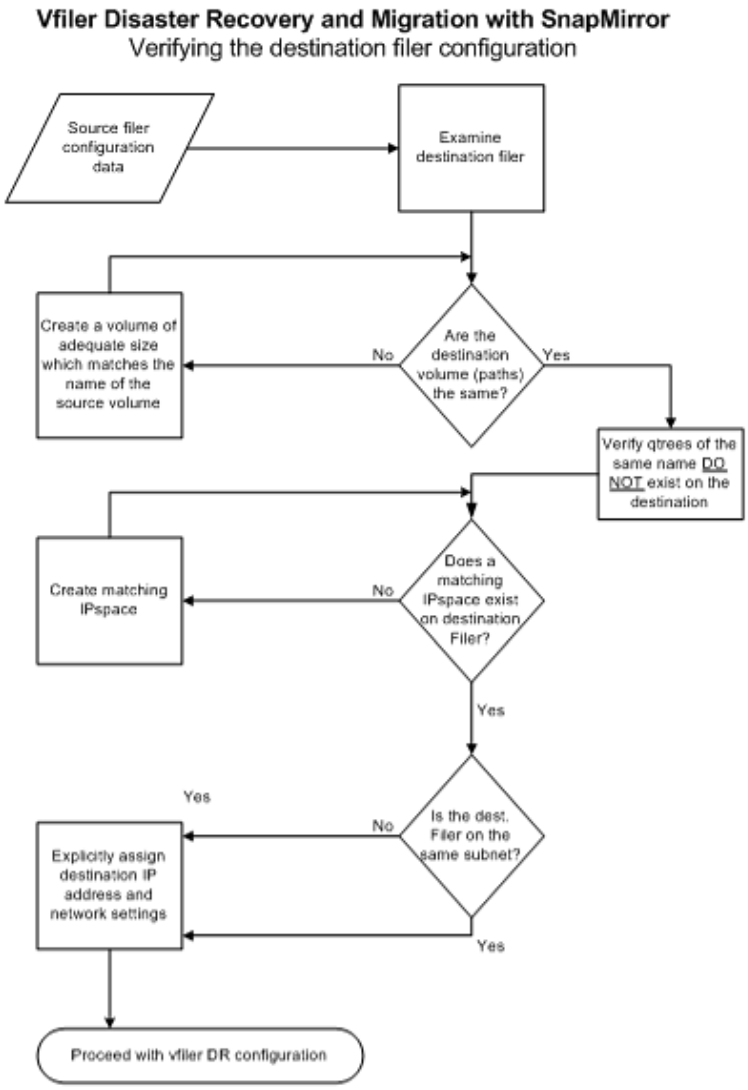

Figure 5: Configuring MultiStore for DR and Migration

Vfiler Migration with SnapMirror

The process of establishing SnapMirror relationships for virtual filers is automated by MultiStore to simplify administration and ensure accuracy. It's still necessary to ensure that the source and destination filers are licensed and appropriately configured, but the MultiStore software takes care of the SnapMirror relationships.

Once a SnapMirror relationship is established, a baseline transfer initializes the mirror to create a replica of the Vfiler on the destination filer. The entire Vfiler including its network and security settings, data, permissions and options are mirrored as one. SnapMirror leverages filer Snapshots to efficiently replicate *incremental* updates. Thus SnapMirror is highly effective and flexible while also efficient in its use of valuable bandwidth.

Each Vfiler's configuration, security, and networking information is stored within its own "root" directory. This allows administrators to mirror or migrate complete virtual filers over the network.

Vfiler Migration with SnapMover

SnapMover is a no-copy data migration solution among filers that share a common storage pool. This can be clustered filers with access to all disks from both nodes, gFilers that can access the same LUNs over Fiber Channel, or a SharedStorage pod with multiple filers connected to a common set of back-end disks. It enables data to be easily and quickly moved from a source node to a destination with no disruption to users.

Data migration is achieved by using standard SCSI-3 semantics to change the ownership of the storage containers (LUNs or individual disk) in the storage pool. The ownership information is updated very quickly and requires little actual I/O, thus making the migration process extremely fast with minimal overhead. Data migration is managed at the volume level (not individual disks/LUNs), making the migration process simple and scalable.

The quick, no-copy migration opens the door to more options for managing storage. It allow IT to easily migrate a data volume from an overloaded filer to another. Volumes can be quickly migrated from a filer that needs to be taken down for maintenance, and just as quickly returned back when the filer is back up. Users are not disrupted during the migration process and continue to access data after migration in the same manner as before. Since the migration is completed in seconds (not hours or days that are typically required by a traditional data replication approach) the IT organization can be much more responsive to its customers and cost effectively provide higher levels of service.

A SnapMover license is required on both the source and destination to enable the data migration capability. It also requires a MultiStore license to provide the transparent data movement benefit. Starting with Data ONTAP 7G, the SnapMover license comes bundled with MultiStore.

2.4) Disaster Recovery Using SnapMirror

Through integration with NetApp SnapMirror, virtual filers can be created and automatically mirrored to other filers over a network for disaster recovery. Integrated mirror relationships are established to automate the creation and synchronization of a *disaster-recovery Vfiler*. In a disaster recovery scenario, the DR Vfiler can be quickly *activated* at one of many possible locations to restore services. Once it is online the mirrored Vfiler is an identical copy of its original source based on the last successful SnapMirror update. The only difference between the original Vfiler and the DR Vfiler is that the latter is inactive until a situation forces it to be activated. Thus the failover can be transparent relative to the applications and users throughout the enterprise.

Migrating a Vfiler vs. Creating a Disaster-Recovery Vfiler

Although the "Vfiler DR" and "Vfiler Migrate" procedures have much in common, they are designed for different purposes.

Vfiler Migrate – Establishes a mirror and replicates the data. Upon completion, the SnapMirror relationship is broken and the source Vfiler destroyed. The source vFiler volumes are not destroyed, so the vFiler can be quickly recreated at the original source if necessary.

*Vfiler DR – Es*tablishes and maintains the mirror relationship and keeps the (active) source and destination virtual filers synchronized. The destination Vfiler remains dormant or inactive until which time it must be activated. If access to the original (source) Vfiler is re-established, the relationship can be effectively "reset" by repairing or replacing the source filer and re-synchronizing the data back to the original source.

Both have a role in most enterprises, and NetApp MultiStore permits both applications with the same license.

Scheduled Replication for DR Virtual Filers

The interval in which SnapMirror updates occur should take into account the amount of data and network speed between source and destinations. For example, if the source and destination filers are connected via a fast LAN, mirror updates can be scheduled to occur at a higher frequency than possible over a slower WAN. For more information on configuration SnapMirror see the Data ONTAP Data Protection Online Backup and Recovery Guide.

In Figure 6, vfiler3 is mirrored from hostfiler1 to hostfiler2. Vfiler3 can be activated on hostfiler2 if hostfiler1 becomes unavailable. Vfiler3 appears on the network exactly as it did before. Users, for example, wouldn't necessarily know a Vfiler had been moved to a different physical filer.

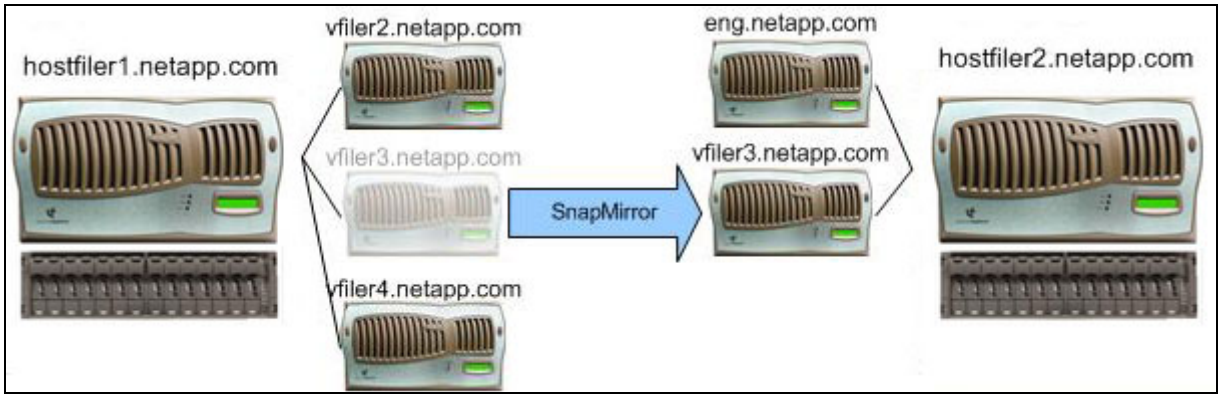

Figure 6: Vfiler Disaster Recovery with SnapMirror

3) NetApp Filers and MultiStore

3.1) Enabling MultiStore Functionality

There is no need to install any additional software in order to take advantage of the MultiStore Vfiler functionality. As long as the filer is running Data ONTAP 6.2 or later, MultiStore can be enabled by entering the correct license code. The license provides instant access to the Vfiler related commands and the configuration process can being immediately.

For information on how to license MultiStore/Vfiler functionality, please see the Data ONTAP 7G MultiStore Management Guide for more information.

Once enabled, the hosting filer becomes also known as vfiler0. Vfiler0 is the name used internally by MultiStore to identify the hosting filer. The hosting filer's actual host name and configuration do not change. Because of the security features of MultiStore, all networking and storage resources need to be owned by one (and only one) Vfiler. Please be aware of the following key behaviors:

By default, vfiler0 owns all of the filer's storage and networking resources.
At any given time, any resources not explicitly assigned to a Vfiler belong to the hosting filer (vfiler0).
Each Vfiler must be assigned at least one storage resource and one networking resource.
Moving, adding, or removing resources between virtual filers only affects the associations between a vfiler and those resources. User data is not affected.
When an administrator creates a Vfiler on a given volume or adds a volume to an existing Vfiler, resources are moved from the hosting filer (vfiler0) to the new Vfiler.
When resources are removed from a Vfiler or a Vfiler itself is removed, the resources associated with that Vfiler are returned to the hosting filer.

The above behaviors help ensure that only authorized users and administrators have access to resources and information in question.

3.2) Storage and Networking Resources

Storage Resources

Storage resources on filers start with a physical file system, or volume. The volume can be either a Traditional Volume or a Flexible Volume. Within a volume, special subdirectories can be created called qtrees that function as logical volumes. Both volumes and qtrees can be assigned to virtual filers as storage resources. Vfiler storage resources (volumes and qtrees) can be added or removed at any time. Removing the storage resource holding the Vfiler configuration information is not allowed without first destroying the Vfiler in question. For a detailed description of volumes and qtrees, please see Chapter 3, “Data Organization Using Volumes and Qtrees” in the *Data ONTAP System Administrator's Guide*.

The filer's entire root volume (typically vol0) cannot be assigned to a Vfiler, though qtrees created within the root volume can be assigned to a Vfiler. All storage resources assigned to virtual filers are secured within the file system. If a Vfiler is deleted, for example, the portion of the file system and data previously owned by that Vfiler may not be accessible to other virtual filers in different security domains.

Volumes or qtrees assigned to virtual filers are known as “storage units.” The first assignment of a volume or qtree to a Vfiler is known as the primary storage unit. The primary storage unit contains an /etc directory that consists of a subset of the files normally found in the hosting filer's /etc directory, specific to that Vfiler's configuration. Examples of this information are UNIX NIS or Windows domain information, configuration options, disk quota assignments, UNIX/Windows user mappings, DNS namespace and server information, etc.

As a result, each Vfiler independently stores information about its membership in Windows and/or UNIX security domains. One Vfiler can serve UNIX clients only, another can be a member of a Windows 2000 or Windows .NET Active Directory Domain, a Windows NT 4.0 and NIS Domain, etc.

Resources assigned to a Vfiler are owned only by that Vfiler. In other words, each Vfiler and its resources are distinct and meaningful only to itself and the security environment (e.g., Windows Domain) to which it belongs. The hosting filer administrator or an administrator responsible for vfilerA does not necessarily have any access to resources assigned to vfilerB.

Networking Resources

There are many choices available when it comes to choosing network cards for a filer. However, the filer's physical network cards (interfaces) can also be logically configured in combinations of up to 128 virtual interfaces, each of which is indistinguishable from a physical interface when it comes to assigning IP addresses.

Virtual filers can be created and assigned IP addresses that correspond to any network interface (logical or physical) on the filer.

Each physical interface can be used individually, and assigned a base IP address

The filer's physical interfaces can be combined into virtual interfaces or VIFs, which are link aggregations or “trunks” of (up to 16) combined links providing fault tolerance and improved throughput

Multiple (logical) Virtual Local Area Network (VLAN) interfaces can be created and associated with any physical interface or VIF

In addition to the flexibility offered by VLANs and VIFs, IP address aliases can be assigned to any interface. IP address aliasing allows more than one IP address at a time to be associated with an interface.

For more information on configuring network interfaces, VIFs and VLANs, please see the Data ONTAP Network Management Guide.

IPspaces

Virtual filers can be grouped into IPspaces, each of which represents a distinct networking environment. An IPspace defines an IP address space that is separate from other IPspaces. Each IPspace maintains its own distinct routing table and no cross-IPspace traffic is routed. In addition to the always present default-ipspace, up to 100 separate IPspaces can be created on a filer.

IPspaces are beneficial in environments where a filer must host storage for two separate departments or organizations, each with private networks. They are invaluable to service providers who can readily create and allocate a brand new virtual filer with its own secure storage, secure administration, and secure routing to customers. Virtual filers allow service providers to deploy and bill for service within minutes.

IP addresses defined within an IPspace are only meaningful within that IPspace. That means that IP addresses are not forced to be unique across IPspaces, and the same IP address can be used in multiple IPspaces for more standardized management. Virtual filers in different IPspaces cannot internally communicate with each other, even though they are running on the same hosting filer. Incoming network traffic on the filer's network interfaces is internally tagged with the IPspace ID to identify the target Vfiler, and virtual filers use the routing table for their IPspace when communicating with other computers.

By default, a single default IPspace exists, named “`default-ipspace`.” All of the filer's interfaces belong to the `default-ipspace` unless they are assigned to a “non-default” IPspace created by the administrator.

Before an IPspace can be used, it must be assigned a network interface. Any of the filer's physical, VLAN, or VIF network interfaces can be assigned to an IPspace. By creating and using logical VLAN interfaces, more IPspaces can be created than physical interfaces available in the filer.

Note in Figure 7 how each Vfiler has the same IP address. This is possible because each Vfiler is in a different IPspace, each of which functions as if on a separate LAN. They function within a network space that has its own network addressing and routing table. Thus, even though they have the same IP address and are hosted on the same filer, there are no address conflicts because they cannot communicate with one another unless they go through an external network router. This kind of practice may simplify planning and administration because all “customers” have similar configurations, with servers which have the same role using the same IP address across customers.

Figure 7: Simplified Management and Security using IPspaces

Using VLANs with IPspaces

It is beyond the scope of this paper to describe advanced networking concepts and VLAN functionality in detail. For the completeness of this paper, a few basic concepts in relation to Ethernet networks and VLANs will be briefly discussed.

Ethernet switches learn about the networks connected to them by analyzing the source address of incoming frames. The network addresses and ports on which they were discovered are stored in tables and used as the basis for traffic forwarding. When a switch receives a network frame on one of its ports, it analyzes the frame to see if the destination matches any of the destinations known on its other ports. If a match is found, the frame is forwarded or "switched" to that port only. If no match is found, the frame is forwarded to all ports.

Unicast frames contain a destination address, while frames with no destination address, such as broadcasts (to all hosts), do not. By default, switches still forward broadcast (and multicast) frames to all ports. Therefore, the mixture and volume of network traffic can still have a profound effect on network performance and reliability.

VLANs can be defined on network switches that confine intra-segment traffic and broadcasts only to VLAN members. Many modern network switches support the creation of port-based or protocol-based VLANs. Port-based VLANs are defined by designating specific ports on a switch as members of a VLAN. Protocol-based VLANs limit traffic to VLAN members according to certain protocol criteria, regardless of which ports they are connected to.

Data ONTAP supports the creation of logical VLAN interfaces that support VLAN tagging. A VLAN tag is a unique identifier that indicates the VLAN to which a frame belongs. VLAN interfaces are assigned IP addresses like any other interfaces in the filer.

When an IPspace is associated with a VLAN interface, virtual filers within that IPspace only communicate with other computers that are members of the same VLAN.

In Figure 8, the filer has two Gigabit Ethernet network interfaces that are combined into a filer VIF (trunk) for maximum performance and reliability. Four VLAN interfaces are defined on the filer's "tr1" VIF interface that correspond to the VLANs defined on a network switch. There are four IPspaces, each of which are assigned to one of the four VLAN interfaces. There is one Vfiler in each IPspace whose IP addresses are associated with each IPspace/VLAN.

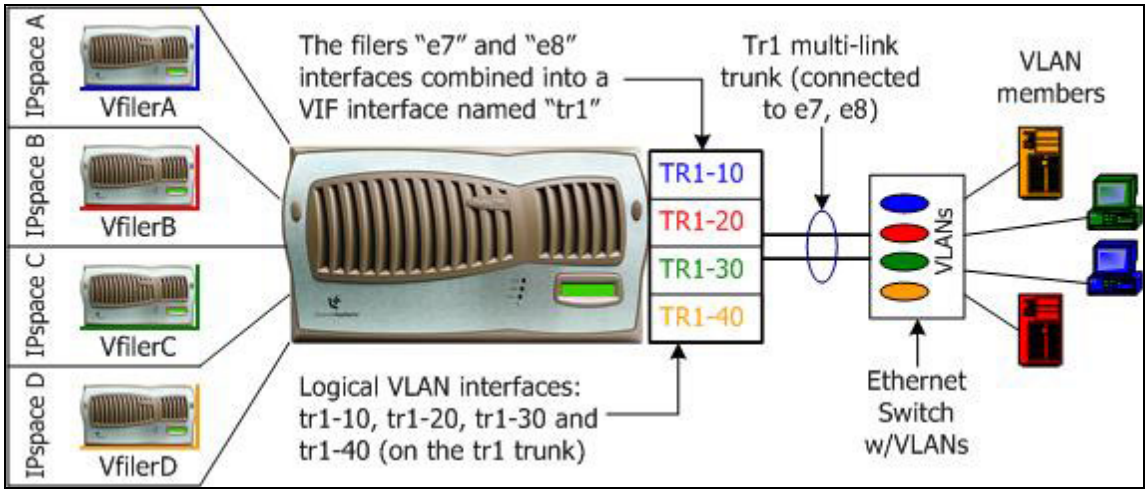

Figure 8: Using Virtual Interfaces, VLANs, and IPspaces with a single network switch

The result is that each Vfiler functions as if it's on its own local area network. It should be noted that each one of the virtual filers in the example could also be members of different security domains. This secure architecture allows for maximum flexibility and allows the filer to determine the target VLAN, IPspace, and Vfiler.

For more information on designing high-availability storage networks, please see the "Storage Networking by Network Appliance and Cisco Systems: High Availability for Network-Attached Storage" white paper. The paper is available online at: http://www.netapp.com/tech_library/3115.html.

3.2) Higher-level Integration

Virus Scanning

Data ONTAP includes integrated antivirus server support for data accessed by Windows computers. The efficient, on-access nature of the NetApp antivirus architecture ensures files are scanned before allowing any reads, writes, or changes to data.

In order for the virus scanning server(s) to intercept file I/O requests, antivirus servers must register with the hosting filer or Vfiler. Virus scanning can be enabled or disabled for each Vfiler, and options can be set independently for each Vfiler. Virtual filers can use scanning servers that are either:

Registered with the hosting filer
Registered with a specific Vfiler

For more information on integrated antivirus scanning for NetApp filers, please see the *Antivirus Scanning Best Practices Guide* available at: http://www.netapp.com/tech_library/3107.html

Quota Management

Disk quotas can be managed on a per-Vfiler basis on the qtrees and volumes owned by a particular Vfiler. However, if a qtree owned by a Vfiler resides in a volume owned by the hosting filer, the hosting filer administrator can also specify a quota for the qtree. The qtree cannot exceed the specified limit in the more restrictive of the two quotas (the lesser quota takes precedence.) Quota settings are preserved when virtual filers are duplicated via SnapMirror to other hosting filers. For a complete list and description of Vfiler-related commands, please see the Data ONTAP MultiStore Management Guide.

Windows Administration

Once all of the necessary setup steps are complete, the Vfiler shows up on the network just like a regular filer. For example, Figure 9 shows how the hosting filer (`sparky`) and its two virtual filers `corpdata` and `engdata01` appear as individual servers within the Active Directory Users and Computers administrative window.

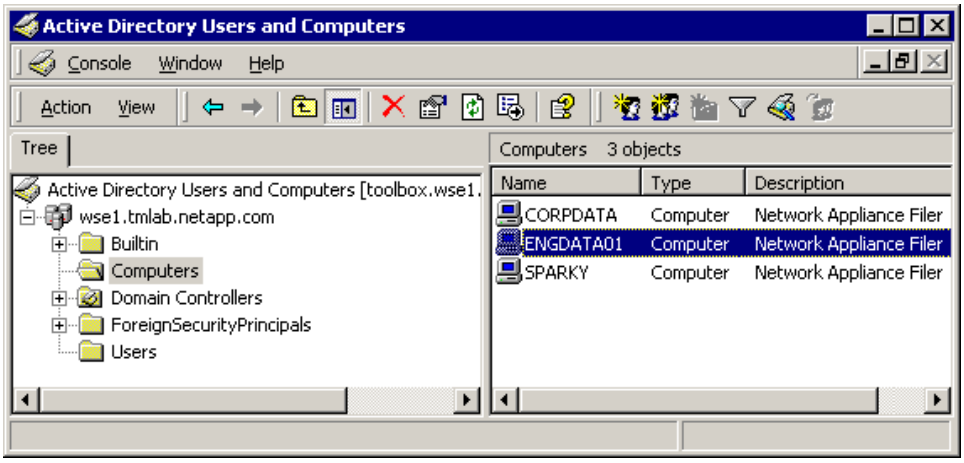

Figure 9: Windows Active Directory sees Virtual Filers like any other filer

Some aspects of filer or Vfiler management can be performed by right-clicking on the filer and selecting the Manage menu option. Figure 10 shows the user sessions connected to `corpdata`. Note how `corpdata` and `engdata01` appear on the network as native active directory servers.

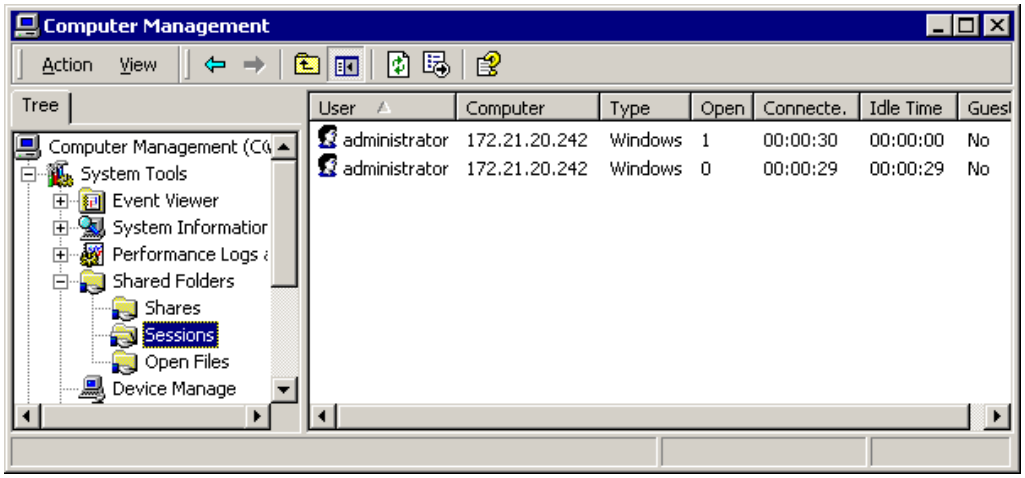

Figure 10: Some Vfiler management can be done directly through Active Directory

4) Summary

NetApp MultiStore allows a filer's storage and networking resources to be partitioned into multiple virtual filers, each of which appears as an individual filer on the network. The dynamic, flexible nature of virtual filers simplifies UNIX and Windows storage consolidation, and offers unified access to data over both CIFS and NFS. Virtual filers can be wholly mirrored to other filers over a network for data migration, disaster recovery, or load balancing purposes. Virtual filers can be grouped into private network address spaces and belong to different security domains for maximum security. MultiStore simplifies data storage and management for companies or service providers requiring highly available storage and secure, multi-domain flexibility. These capabilities lead directly to added value for organizations using MultiStore and SnapMover:

Different organizations may be responsible for managing the hosting filer and/or each Vfiler. Organizations benefit by centralizing management of the hosting filer, while allowing client departments to independently, securely, and responsively manage their own data. The result is **efficiency and great service**, with the some functions reliably centralized and scaled, while others are distributed closer to the end user who ultimately benefits from responsive and more tailored service.

Vfiler administrators for different organizations do not have permission to access other virtual filers unless explicitly allowed. This allows for more **flexible management policies and greater security**.

Virtual filers may be managed using the Command Line Interface (CLI), the web-based NetApp FilerView application, the centralized NetApp DataFabric Manager, or the Manage ONTAP API. Many management options means it is possible to integrate filers and MultiStore into a wide range of existing operations and processes. The result is **rapid deployment and reduced process re-engineering costs**.

Virtual filers can be easily created and destroyed as needed to accommodate temporary projects or changing project priorities. No additional hardware needs to be deployed or managed. The **unprecedented ease and speed opens doors to more efficient operations and project opportunities**.

As business needs and workloads change, virtual filers may be migrated to ensure that critical projects receive the required level of service and response times, while retired projects are still available online via lower-tier storage. As priorities change, the administrator simply adjusts Vfiler deployments. This flexibility leads to **greater return on investment (ROI) and lower total cost of ownership (TCO).**

Storage management is virtualized, allowing the IT architect to separate many of the logical needs of the infrastructure from the physical hardware used to build it. Because less hardware is required, **management and acquisition costs can be reduced, and adaptability improved**.

In conjunction with SnapMover technology, a Vfiler can be migrated to another NetApp filer in seconds, with no data being copied and no interruptions to the user. Combined with Grid technology, the load balancing capabilities provide **opportunities for unprecedented application scaling efficiencies**.

MultiStore and SnapMover technologies can also be used with Vfiler migration for system maintenance with no operational interruptions. Non-disruptive maintenance delivers **higher levels of service and allows more profitable operations** by being able to offer more stringent Service Level Agreements.

SnapMirror can be used to replicate virtual filers to one or more target filers, where the mirrored virtual filers can be quickly activated for disaster recovery purposes. This significantly **simplifies disaster preparedness, and reduces the chances of anything else going wrong** in the stressful minutes following a disaster.

NetApp®

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