Storage Virtualization and DR Using MultiStore® (vFiler™)

Kiran Sreenivasamurthy, Network Appliance, March 2006, TR 3462

Executive Summary

This technical report describes storage virtualization using MultiStore and its benefits. It provides an architectural overview and configuration best practices. The report also highlights the differences of configuring MultiStore for either disaster recovery or migration. This report provides issues to be aware of while configuring disaster recovery and migration as well as an overview of the new features that will be available in Data ONTAP® 7.2.

The report provides sample usage scenarios which can be used as a benchmark to configure virtual storage controllers in a new environment.

Table of Contents

1. Introduction

The onset of the digital age and numerous software applications, e-business, and other enterprise technology trends such as ERP, CRM, and data warehousing are doubling the amount of corporate data every six months. Businesses are looking for technologies to help them manage and store the information flowing in and out of their computer systems at an alarming rate, 12 times the data stored in 1998.

MultiStore, also known as vFiler, is an optional software solution that enables secure, multiprotocol storage consolidation across enterprises. It provides secure partitioning of network and storage resources and enables storage consolidation for multidomain and multiserver configurations. In addition, it reduces management cost by reducing the number of storage systems that the storage administrators have to administer, thus reducing the total cost of ownership.

2. vFiler Benefits

MultiStore technology, developed by NetApp, enables the physical resources of a storage system to be logically partitioned to form separate virtual storage controllers. These virtual storage controllers are also referred to as “vFiler” controllers. Logically separating the resources of a physical storage system provides the following benefits:

- **Virtualization**

Virtualization provides a layer of abstraction, decoupling the physical resources like CPU and system memory of the physical storage system. It provides a logical view of both storage and computing resources. Virtualization hides the complexity and simplifies storage and system management.

- **Consolidation and Ease of Management**

MultiStore technology provides an efficient architecture for consolidating multiple physical storage systems into a smaller number of systems. From the end user’s perspective, each virtual storage controller appears as a separate physical storage system with a unique IP address. For example, storage for multiple business units can be consolidated into a smaller number of systems, increasing the return on investment and also easing management challenges. Application service providers can also consolidate the storage needs of their customers.

- **Security**

Security is one of the key concerns when storage is consolidated either within an organization or by an application service provider. A virtual storage controller provides a confined environment. The data owned by a virtual storage controller cannot be accessed by any other virtual storage controllers even though they are hosted on the same physical storage system. All requests for data access owned by a virtual storage controller are tagged with its context, making it impossible for unauthorized access to data.

- **Delegation of Management**

MultiStore provides the ability to delegate management access based on their role. The virtual storage controller administrators can have different access rights compared to physical system storage administrators. This provides another layer of security. For example, in an application service provider model, a single storage system hosts the data owned by multiple customers. Controlling access rights and delegation of storage management enables data security and eases storage management.

- **Disaster Recovery**

Disaster recovery is essential to protect business-critical data and mitigate the risk of a catastrophic data availability failure. MultiStore provides an easy to deploy and manage disaster recovery solution that improves recovery time and lowers management costs. The use of virtualization technology removes the requirement that the primary and the backup systems be identical. In case of a disaster, MultiStore technology provides the tools needed to restore access to clients without having to make any client configuration changes.

- **Workload Management**

MultiStore technology provides an efficient mechanism to perform workload management by migrating virtual storage controllers across storage systems. Upgrading hardware or decommissioning old hardware can be accomplished with limited disruption of data access and zero configuration changes on the clients.

3. Best Practices and Deployment

3.1 Introduction to Qtrees and Traditional and Flexible Volumes

Traditional Volume and Flexible Volume

A traditional volume is a collection of physical disk space (4K blocks) whose entire capacity constitutes a single volume. These disks are arranged in RAID groups depending on the RAID level. The size of the traditional volumes can be dynamically increased but not reduced.

FlexVol™ technology, introduced with Data ONTAP 7G, increases flexibility and storage efficiency through the use of flexible volumes, which decouple the direct connection between volumes and their associated physical disks. Flexible volumes are linked logically to the underlying physical storage through aggregates, which isolate the flexible volumes from the physical disk's infrastructure. An aggregate can host multiple flexible volumes. A flexible volume can be dynamically increased or decreased.

Qtree

A qtree is similar in concept to a partition on a Windows® or a UNIX® host system. A volume can have multiple qtrees. Quotas can be applied to a qtree to limit its size. As a special case, a qtree can be the entire volume. A qtree is more flexible than a partition because the size of a qtree can be changed without any downtime to the users.

3.2 Virtual Storage Controller

3.2.1 Introduction

A virtual storage controller is a lightweight instance of a multiprotocol server. Physical resources of the storage system such as system memory and CPU are shared between virtual storage controllers. A virtual storage controller consists of data stored in a volume or a qtree, the IP address(es) necessary to reach the virtual storage controller, and the security and other attributes associated with the data. From the client systems and management software perspective the data is completely secured and isolated from all other virtual storage controllers.

The conceptual difference between a virtual storage controller and a virtual machine is that a virtual machine (VM) in the traditional sense is a partition resulting from static, hardware-assisted partitioning of physical hardware resources (including storage disks and network interfaces) of a computer system. A VM can run an independent copy of the software that can normally run on the host computer system, and can thus appear as an independent computer system. A virtual storage controller, in contrast, is a dynamic partition associated with software objects in Data ONTAP that represent various storage and networking capabilities. Therefore, volumes and qtrees are assigned to virtual storage controllers and not physical disks. Also, we assign IP addresses to virtual storage controllers and not network interface cards. A virtual storage controller is thus differentiated from a VM in that multiple virtual storage controllers can share physical hardware resources. Furthermore, the partitioning of software objects representing storage and networking capabilities into a virtual storage controller is much more flexible and dynamic than the hardware-based static partitioning of a physical machine into VMs and makes them significantly easier to manage than a VM.

A virtual storage controller's configuration allows it to store and retrieve data in the correct context in its storage units. This information also allows the virtual storage controller to correctly interpret the access-control and security-related metainformation embedded in its storage.

The dynamic association of a virtual storage controller with its storage and networking resources makes the movement of resources a relatively easy operation.

Enabling a MultiStore license will create a virtual storage controller named “vFiler0.” “vFiler0” is referred to as a default virtual storage controller and cannot be renamed. Any other virtual storage controller created on the storage system is referred to as a nondefault virtual storage controller.

3.2.2 Deployment

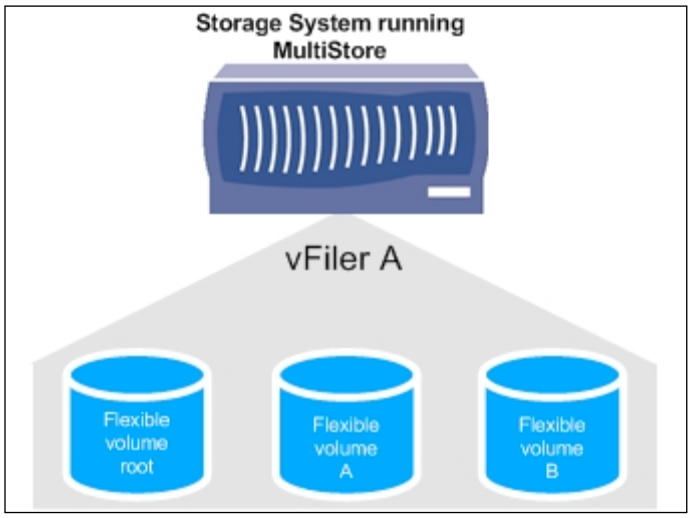

Figure 1) Virtual storage controller configuration.

Figure 1 shows a virtual storage controller configuration within a physical storage system running MultiStore.

A virtual storage controller configuration includes:

- Quota, exports, and log information
- Configuration for hosts, DNS, and NIS
- CIFS domain info, local users, groups, shares
- The subset of storage system options that are specific to a virtual storage controller
- Virtual storage controller registry data

The storage unit containing the entire virtual storage controller’s configuration is permanently associated with the virtual storage controller. This storage unit is also known as the primary storage unit. This storage unit can be disassociated only if the virtual storage controller is destroyed. The other resources, such as data volumes or qtrees or the network interface, can be added or deleted from the virtual storage controller. The deleted resources will be owned by the default virtual storage controller. These resources must be manually brought offline. The configuration information controls how the data is being served—for example, using NFS over UDP or NFS over TCP.

Points to keep in mind when deciding on the storage unit for the virtual storage controller:

- It is advantageous to have the primary storage unit be its own flexible volume and the data accessed stored in separate storage units. Data volumes can be reassigned to different virtual storage controllers without copying the data. This will enable you to manage the workload between them. Multiple virtual storage controllers can share the same aggregate.

- Though the primary storage unit can be a qtree or a traditional volume, we recommend that the primary storage unit be a flexible volume.

The maximum (hard limit) number of virtual storage controllers that can be created depends on the amount of physical memory and the Data ONTAP version running on the system. On a production environment, the best practice is to understand the workload on the storage system. Tables 1 and 2 list the maximum number of virtual storage controllers that can be created, including vFiler0.

Data ONTAP Version 7.0.3	
STORAGE SYSTEM	**HARD LIMIT ON THE NUMBER OF VIRTUAL STORAGE CONTROLLERS**
FAS250, FAS270	11
FAS3020	26
FAS960, FAS980, FAS3050, R200	33

Table 1

Data ONTAP Version 7.2	
STORAGE SYSTEM	**HARD LIMIT ON THE NUMBER OF VIRTUAL STORAGE CONTROLLERS**
FAS250, FAS270	11
FAS3020	26
FAS960, FAS980, FAS3050, FAS6030, FAS6070, R200	65

Table 2

In a clustered environment virtual storage controllers can be configured up to the hard limit, as mentioned in Tables 1 and 2 on each node. Note that in the case of a system failure the surviving head will host the entire partner's virtual storage controller. In effect the surviving head will host double the hard limit, and this can lead to bottlenecks in the functioning head's hardware resources. For example, configuring additional network interface ports to be used by the partner in case of a failover can reduce the network bottleneck. Things to keep in mind when configuring virtual storage controllers:

- A physical storage system supports a predefined maximum number as per Tables 1 and 2.
- FCP is not supported.
- The following protocols are supported: NFS, CIFS, iSCSI, HTTP, NDMP, and FTP.

- Login shell and console are not available; you have to switch context from vFiler0 or execute commands remotely.
- A subset of Data ONTAP commands is available.

3.3 Virtual Storage Controller DR and Migration

DR (disaster recovery) and migration use SnapMirror® as the base technology for data replication. Virtual storage controller DR and migration in asynchronous mode are supported on all storage systems. The storage units when configuring a DR or migration can be a volume or a qtree.

If the storage unit is a volume, then the corresponding source and destination volumes owned by the virtual storage controller must be of the same type (either traditional or flexible volume). For example, if the source has two flexible volumes (flex1 and flex2) and a traditional volume (trad1) as its resource, then the destination must also have two flexible volumes and one traditional volume (flex1, flex2, and trad1).
If the storage unit is a qtree, then the qtree can reside in either a flexible volume or a traditional volume. For example, if the source has a qtree (tree1) on a traditional volume, then the destination can have the qtree (tree1) on a flexible volume.

3.3.1 SnapMirror

SnapMirror provides a simple, efficient way to replicate data between storage controllers using multitransport (IP and FC) and frequency choice (sync, semi-sync, and async) while only transferring changed data. Based on the customer needs and resource availability, storage administrators can choose either or both technologies.

The total number of concurrent SnapMirror streams is based on a particular platform, and it ranges from four to 128. For FAS storage controllers with ATA disks this number is reduced by half. The total number of SnapMirror streams depends on the type of storage system and is not equally divided among the virtual storage controllers. The number of SnapMirror streams limits the number of DRs or migrations that can concurrently run on the entire physical storage system. SnapMirror configurations are stored within the virtual storage controller's root directory. If there are multiple virtual storage controller administrators on the physical storage system, a good communication protocol has to be established so that they do not impact each other. The exact number of SnapMirror streams for a particular platform can be found on NOW™ (NetApp on the Web).

If more than the allowed number of asynchronous SnapMirror volumes or qtree replications is scheduled to run at the same time, each additional transfer will generate an error message stating that resource limits have been reached. Each transfer beyond the limit is retried every minute until it succeeds, SnapMirror is turned off, or the update is terminated.

Synchronous SnapMirror and semi-synchronous transfers work only at the volume level. Asynchronous transfers can work on both volumes and qtrees. Synchronous SnapMirror replicates data from the source to the destination at the same time that data is written to the source volume. In case of synchronous SnapMirror, the client gets an acknowledgement only after the data is written on both nodes. Synchronous SnapMirror transfers are more expensive than asynchronous SnapMirror transfers. Refer to the Synchronous SnapMirror Design and Implementation Guide for more information.

3.3.2 Virtual Storage Controller DR

Disaster recovery solution restores access data quickly, reducing user downtime. In a DR configuration, the source system remains active serving data to its clients, and the destination system remains inactive but ready to be activated in case of a disaster. It is recommended to have the disaster recovery site geographically farther from the source to recover from any sitewide disaster. The activation process has to be performed manually.

DR is configured on the destination storage system and has three logical steps:

1. Configure: Configure initiates the mirroring of the remote virtual storage controller to the local virtual storage controller using SnapMirror.

2. Monitor: Monitor the status and success of the SnapMirror transfer.

3. Activate: In the event of a disaster, activate the remote system on the local machine.

Configuring a DR creates an entry in the sampmirror.conf file on the destination system with the default schedule. Since there is an upper bound on the number of concurrent SnapMirror transfers, it is recommended to stagger the transfers by editing the entries in the snapmirror.conf file. The frequency for the SnapMirror transfers depends on the:

- Number of SnapMirror transfers on the source and destination storage systems
- Duration of SnapMirror transfers
- Rate of change of data on the source storage system
- Service-level agreement

Advantages and Disadvantages of DR

VIRTUAL STORAGE CONTROLLER DR	
ADVANTAGES	**DISADVANTAGES**
Provides data redundancy	Dedicated hardware required
Efficient use of network bandwidth since only delta changes are copied	
Protects data from site disaster	
No client configuration changes are required after activating the DR site if we choose the same IP address for the virtual storage controller on the destination machine	
Access to data can be easily and quickly restored. The data is kept in sync using the SnapMirror technology.	

Table 3

Virtual storage controllers have the capability to resynchronize the changes on the disaster site back to the source storage system. If the secondary storage system has to be activated (made the primary) because of a disaster or internal auditing, it can be reverted using the “resync” functionality. The “resync” feature copies the delta changes from the secondary system back to the primary system.

3.3.3 Virtual Storage Controller Migration

Migration moves the specified virtual storage controller from the remote storage system to the local storage system. Migration is initiated on the destination storage system that will host the virtual storage controller after the migration. Migration across storage systems enables workload management. Migration automatically destroys the source virtual storage controller and activates the destination, which starts serving data to its clients automatically. Only the configuration is destroyed on the source, not the data. The migration process takes more time than activating the DR destination site since it has to do a level 0 copy of

the data. Migration can also be used to perform hardware maintenance on the storage systems with limited downtime.

This operation has three logical stages:

1. Start: Level 0 dump of data is performed using SnapMirror technology. Depending on the amount of data, this step can take a long time. During this phase the clients still have access to the source storage system, and the destination storage system is not yet activated.
2. Monitor: Monitor the status and success of the data migration.
3. Complete: The virtual storage controller on the destination is stopped to prevent data access, and an incremental copy of the data is performed. This limits the downtime for data access. Once the data migration is complete, the parameters are set on the destination storage system and activated. The virtual storage controller on the source system is then automatically destroyed. Only the configuration is destroyed, not the data.

Advantages and Disadvantages of Virtual Storage Controller Migration

VIRTUAL STORAGE CONTROLLER MIGRATION	
ADVANTAGES	**DISADVANTAGES**
Upgrade or decommission old hardware	Migration might take a long time depending on the amount of data
Dedicated hardware is not required	Level 0 copy each time there is a data migration
Balance workload by migrating virtual storage controller off the busy storage controller	Heavy use of network bandwidth during migration
No client configuration changes are required after the migration if we choose the same IP address for the virtual storage controller on the destination machine	
Data consistency is preserved by stopping the virtual storage controller on the source system and then updating the data on the destination system using SnapMirror updates	

Table 4

The limit on the number of concurrent SnapMirror transfers limits the number of migrations that can be executed simultaneously.

3.4 IP Space

An IP space is a logical qualifier for IP addresses which allows a single storage system to properly handle possibly overlapping private IP addresses. Each interface belongs to only one IP space, but an IP space can have multiple interfaces.
DR or migration on the destination site can be configured with the same or different network configurations. The source and the destination virtual storage controllers should be in the same IP space. The destination virtual storage controller can have a different IP address, a different netmask, and a different default gateway. As long as there is network connectivity between these default gateways, the clients can still access the resources.

The disadvantages when the network settings differ are:

- Clients have to be reconfigured if they are using the server's IP address to access the resources.
- If the clients are accessing the resources using the server's name, then the DNS server has to be updated with the new IP information. If you are creating an alias in the DNS, the source storage system should be down or the application should support load-balancing architecture. This will avoid data loss and/or data inconsistency.

Whether the IP addresses are the same or different, the IP address must be in the same IP space. The DR activation process is not automatic. A well-defined procedure has to be developed to handle the DR activation procedure. This will help reduce the downtime and provide higher data availability.

3.5 New Features of MultiStore in Data ONTAP 7.2

3.5.1 SnapMirror and SnapVault® Integration

SnapMirror Integration

SnapMirror integration provides the ability to manage SnapMirror relationships within a virtual storage controller. Virtual storage controllers can own a volume(s) or a qtree(s), which are the sources or destinations of SnapMirror relationships. The SnapMirror relationship is maintained only when the virtual storage controller is migrated or failed over if the volume or the qtree is the source of the SnapMirror relationship. If the volume or the qtree is the destination, then the SnapMirror relationship will be broken off. As illustrated in Figure 2, when "vFiler A" hosted on storage system SS1 was migrated over to storage system SS3, the existing SnapMirror relationship on a volume owned by "vFiler A" was also migrated over to storage system SS3. The number of SnapMirror streams remains the same as with the previous versions of Data ONTAP. The best practices for configuring the number of SnapMirror transfers are still relevant. The factors that affect the frequency of SnapMirror transfers still apply.

Things to keep in mind when configuring SnapMirror on nondefault virtual storage controllers:

- If a virtual storage controller owns a qtree, but not the hosting volume, QSM on this qtree is not supported.
- If a virtual storage controller is rooted on a qtree, QSM is not supported.
- SNMP is not supported.
- Since tape drives are not supported on nondefault virtual storage controllers, the following commands are not available:
  - SnapMirror store or retrieve
  - SnapMirror use

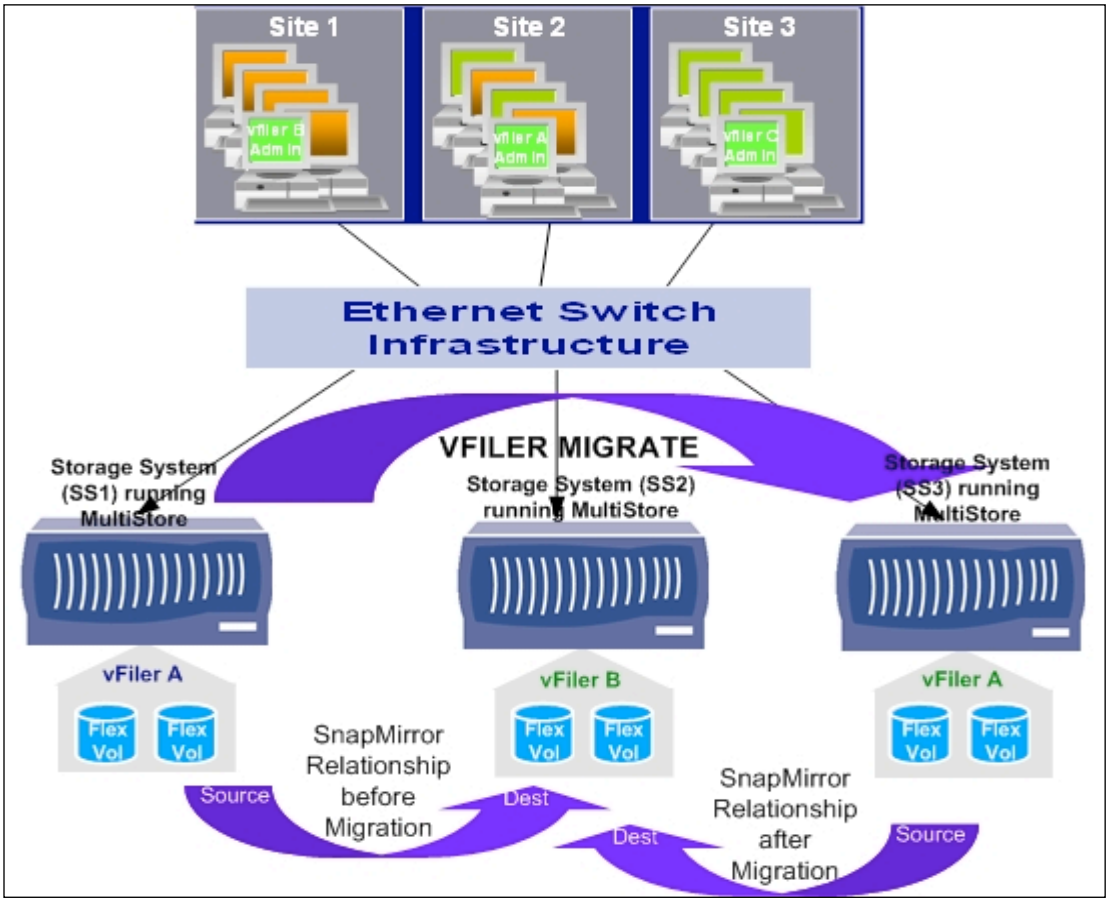

Figure 2) SnapMirror integration with virtual storage controller.

Figure 2 illustrates the SnapMirror relationship being maintained after migration.

SnapVault Integration

SnapVault integration will help maintain the relationship across virtual storage controller migration. Enabling and controlling access to SnapVault can be done per virtual storage controller. The destination can be updated from both the physical storage system and the virtual storage controller.

3.5.2 DataFabric® Manager Integration

Data ONTAP 7.2 enables virtual storage controller discovery, monitoring, and quota management from clients running DataFabric Manager (DFM) version 3.3 and above. DFM uses SNMP queries to the physical storage system to monitor and discover virtual storage controllers. Using DFM, storage administrators will be able to manage quotas. DFM provides the functionality to manage virtual storage controllers from a physical storage system. The virtual storage controller management functions supported are:

- Creation/deletion
- Migration
- Resources assignment (storage and IP addresses)
- Disaster recovery
- Resync in DR relationships

Certain operations—for example, using SnapMirror and Snapshot™ and block-level management through DFM—are not available at the virtual storage controller level.

3.5.3 SnapDrive® Integration

SnapDrive is an enterprise-class storage and data management solution available for both Windows and UNIX platforms. SnapDrive is installed on either the Windows or the UNIX host systems to simplify block-based access protocol management. Data ONTAP 7.2 allows SnapDrive to support virtual storage controllers. SnapDrive integration will enable creating new LUNs, connecting to existing LUNs, configuring and managing iSCSI connections, and managing Snapshot copies. Since FCP does not support virtual storage controllers, SnapDrive support for managing FCP connections is not available. SnapDrive operations on virtual storage controllers that do not own the volumes are not supported. This will prevent virtual storage controllers rooted on qtrees on the same volume from competing for resources and/or interfering with operations.
There is no change to the existing user interface. The changes made for the integration are backward compatible with versions greater than SnapDrive version 3.1.

3.5.4 NDMP Integration

The ndmp service has been integrated into virtual storage controllers in the 7.2 release of Data ONTAP. Both ndmpd and ndmpcopy commands can be executed within a virtual storage controller context.
The storage administrators can perform ndmpcopy at their own schedule. The overhead of integrating ndmp services is very minimal.
Since the overhead of virtual storage controllers is minimal, the performance of running ndmpcopy within a virtual storage controller is similar to running multiple ndmpcopy sessions on the storage system.

The interface for running ndmpcopy is the same as that of a physical storage system. There are a few limitations when running ndmp service within a virtual storage controller. Please refer to the man pages of ndmpd for these limitations.

One important observation is that the performance degradation is not linear as we add more virtual storage controllers running ndmpcopy. Based on the type of storage system it reaches a plateau and then degrades rapidly.

Things to keep in mind:

- SNMP support for NDMP is not available on virtual storage controllers.
- Local NDMP backups/restores are not supported on virtual storage controllers.
- Restartable dump will not be supported.

3.5.5 SSH Integration

SSH integration will provide customers a secure method of remotely executing commands without logging onto a virtual storage controller. SSH can be used to execute commands remotely but not to log in to virtual storage controllers. SSH supports the same protocols on the virtual storage controller as the physical storage system. Note that there is no SSL support. Since the management of a public/private key pair is enabled in a virtual storage controller context, each virtual storage controller has its own key pair. Currently, 12 simultaneous SSH connections are supported on the storage system. With the integration of SSH the 12 simultaneous connections are shared between the virtual storage controllers. For more information on SSH, please refer to the "Best Practices for Secure Configuration" guide.

3.5.6 Manage ONTAP™ API Integration

Manage ONTAP API is an XML-based interface that is used by various applications developed by NetApp and third-party GUI-based system management tools. Clients invoke Manage ONTAP API interfaces that are handled and serviced by Data ONTAP. With the integration of Manage ONTAP API, Manage ONTAP APIs will be available to the clients in the context of the virtual storage controller. The Manage ONTAP APIs that will be available belong to CIFS, NFS, qtree, and quota groups. Please refer to the user manuals for the complete list of Manage ONTAP API calls that are available.

4. Security

A virtual storage controller is a logical partitioning of the hosting storage system's physical resources. Resources owned by a virtual storage controller cannot be accessed or discovered by other virtual storage controllers. Since data and configurations are encapsulated within a virtual storage controller, it provides a convenient way to implement data migration and disaster recovery.
Multiple storage units (volume or qtree) or networking resources can be added to a single virtual storage controller. If multiple qtrees are configured on a volume and these qtrees are added as resources to multiple virtual storage controllers, then a volume might be shared by multiple virtual storage controllers.

Virtual storage controllers can be assigned their own unique IP address and can either share or have their own physical network interface. Multiple interfaces can be assigned to an IP space, so the sending and receiving of data are confined within this IP space. Each interface can be assigned to only one IP space. Based on the security needs, there has to be a tradeoff between the number of physical network interfaces and distinct IP spaces. Each virtual storage controller maintains its own routing table. It is important to review the routing table configurations to make sure there are no static routes that would jeopardize network security.

When a storage system receives a request over the network, the network driver passes the request to the IP protocol stack. This request is assigned to a context based on the destination IP address and the IP space associated with the network interface. This context is carried with each request throughout the handling of the request. Virtual storage controllers also have their own protocol stack enabling them to listen on their own ports. Since context is carried throughout the request, the same port number can exist in multiple virtual storage controllers.

Similarly, a storage unit owned by one virtual storage controller cannot be accessed by another virtual storage controller. The storage system maintains a table that allows the volume and/or qtree for a file or directory to be mapped to an owning virtual storage controller. The context that is assigned with every request should match the entry in the table for the file or directory that is to be modified. If there is a mismatch, the request fails immediately, preventing access to data that does not belong to the intended virtual storage controller. If a symbolic link resolves to a path outside the virtual storage controller's boundary, the data access will fail since there is a mismatch in the context of the request and the data maintained in the table.

A virtual storage controller storage administrator's access rights can be different than the physical storage system administrator's access rights. Since virtual storage controllers maintain their own password file, separate sets of users can be created. The best practice is to create users within a centralized lookup service like NIS and use role-based access control (RBAC) to limit their access rights. RBAC can be used to limit the set of commands a user can execute.

5. Usage Scenarios

MultiStore is typically used in the following scenarios:

- **Delegation of Management and Security**
  As the data requirement in an organization grows, managing the data is no longer a trivial task. The need for the delegation of data management to multiple administrators becomes more important. MultiStore can accomplish this goal with minimum changes to the existing storage system. The storage system should have free space to host the virtual storage controller's root directory, which is also the primary storage unit. A MultiStore license has to be installed on the storage unit. The additional steps that have to be performed are:
    - Create virtual storage controllers with their own flexible volumes, which are the primary storage units for hosting the root directory.
    - Identify the existing volumes and/or qtrees that each virtual storage controller owns.
    - Add volumes and qtrees as resources to the appropriate virtual storage controllers.
    - Delegate management of these virtual storage controllers to different administrators.

- **Basic Disaster Recovery**
  Basic disaster recovery can be configured over either LAN or WAN. DR is preferably configured over WAN to accommodate site failures. In this configuration, the destination or the secondary site acts as both a backup site and a disaster recovery site. The network traffic over the WAN is directly proportional to the rate of change of data on the primary or the source storage system. If the SnapMirror transfers are scheduled at a low frequency, the loss of data will be large if there is a major failure on the primary or the source storage system.
  Synchronous SnapMirror will add more overhead on both the network and CPU resources, especially if the DR is configured over the WAN.

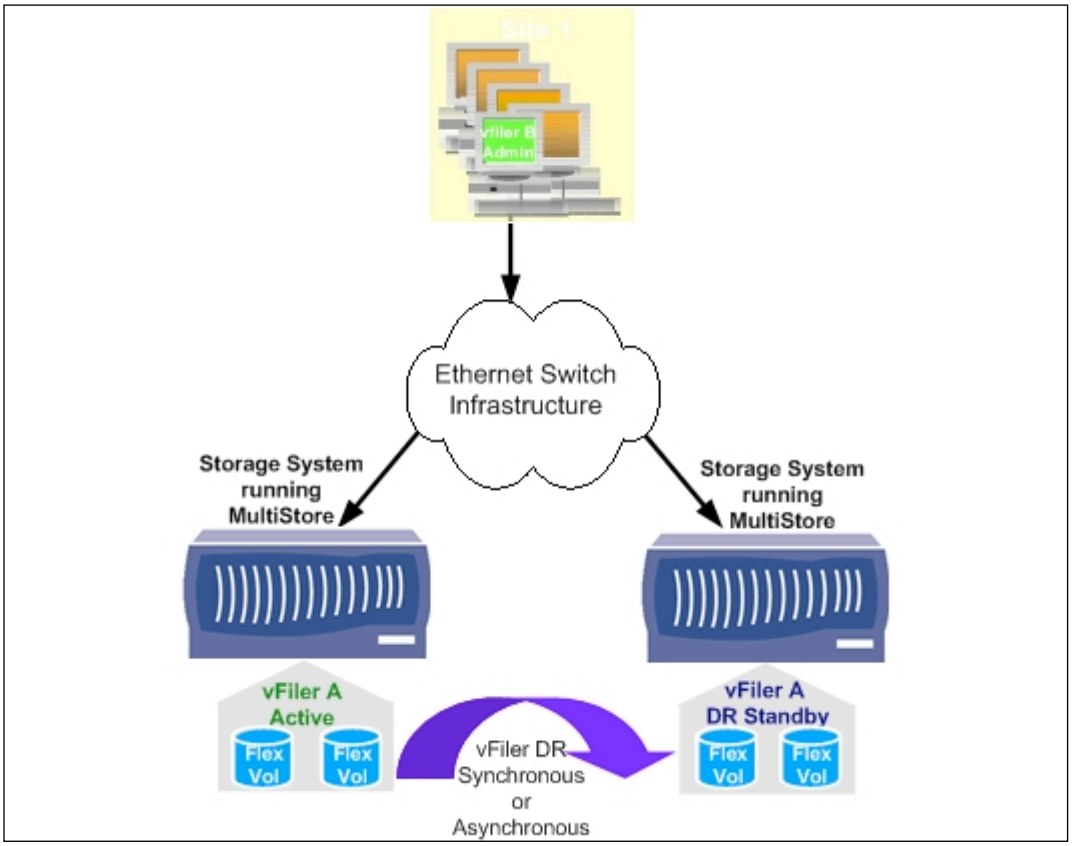

Figure 3) Disaster recovery configuration.

- **Disaster Recovery and Backup**
  This scenario adds an additional layer of data protection compared to the basic DR configuration. In this configuration data is backed up using SnapMirror in synchronous mode at the primary site over LAN. DR can be configured in asynchronous mode over WAN to the DR site, better utilizing network and CPU resources. This architecture still has a single point of failure on the DR site.

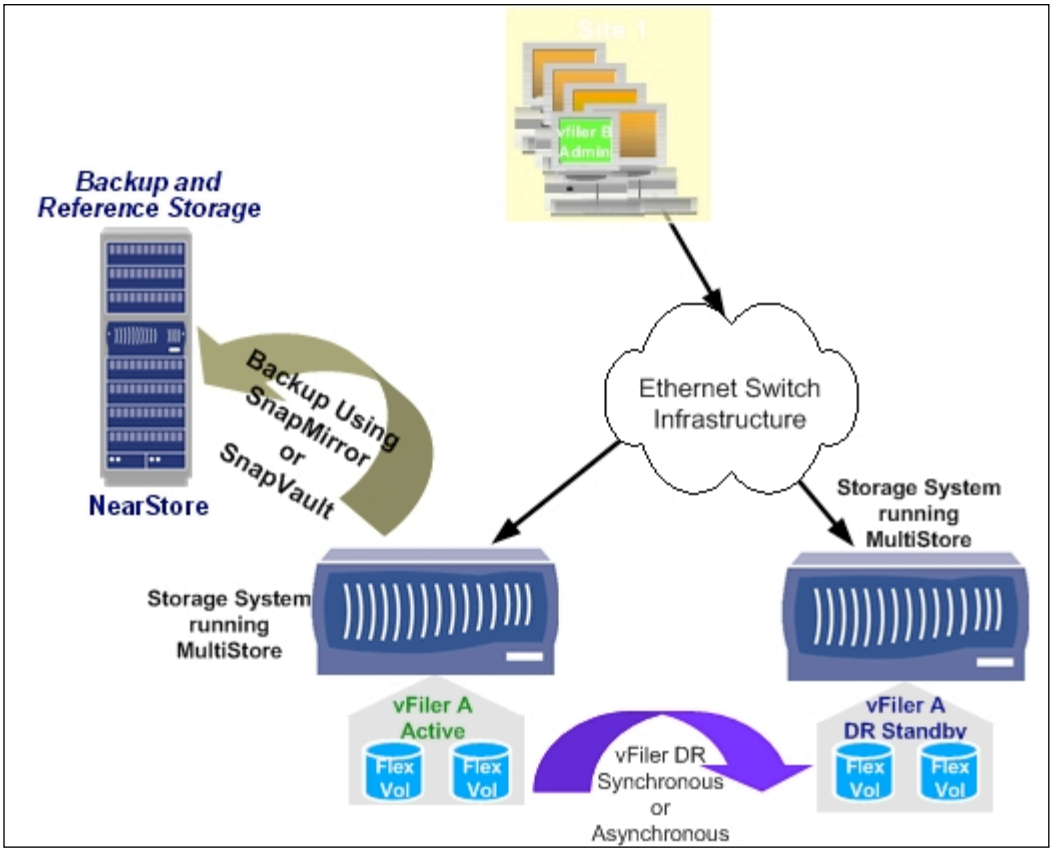

Figure 4) Disaster recovery configuration with a single backup.

- **Disaster Recovery and Two Backups**
  In addition to the advantages of having a single backup, this scenario provides an additional layer of data protection by backing up the data to secondary storage at the DR location. This architecture eliminates the single point of failure of the storage system on the DR site.
  It is important to note that if DR is configured in asynchronous mode, data can be copied to the NearStore® system on the DR site using SnapMirror only in asynchronous mode. However, if DR is configured in synchronous mode, data can be copied to the NearStore system on the DR site using SnapMirror in either synchronous or asynchronous mode.

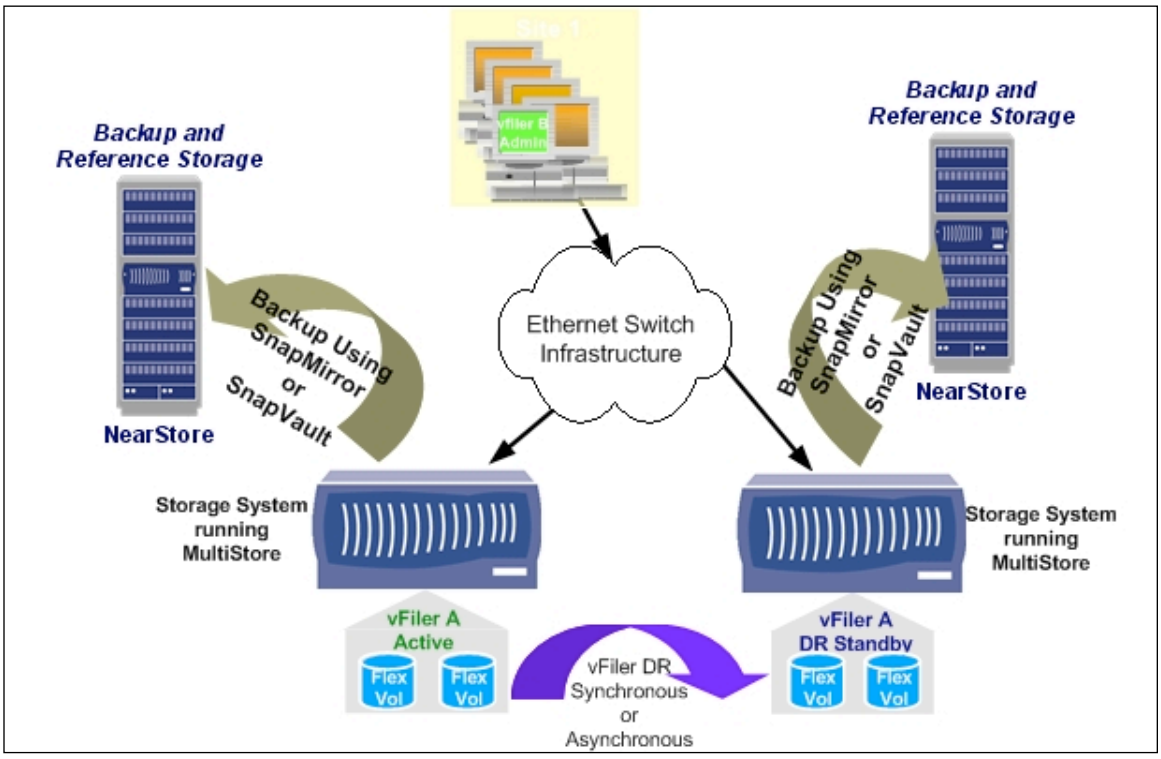

Figure 5) Disaster recovery configuration with two backups.

The cost and data availability are directly proportional to each other. As the requirement for data availability increases, the cost for providing the solution also increases. Figure 6 is an illustration of this relationship. The figure is not drawn to scale and is not an indication of the proportion with which the cost increases as the need for data availability increases.

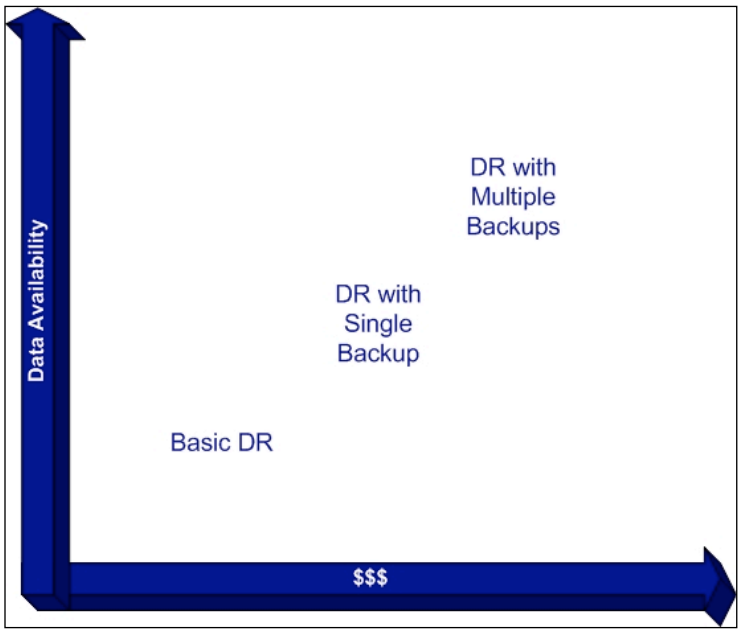

Figure 6) Relationship between data availability and cost.

6. Summary

MultiStore provides a simple, elegant, and flexible architecture to enable storage consolidation, storage virtualization, and service-level model deployment. Storage consolidation and virtualization provide ease of data management, simple disaster recovery configuration and management, and hardware upgrade with minimum disruption. Virtualization also helps in delegating administration with granular access control using RBAC. All these advantages help the user manage and comply effectively with service-level agreements. MultiStore helps the storage administrator to work smarter rather than work harder.

7. Appendices

7.1 References

1. NetApp MultiStore and SnapMover®
   This document describes how NetApp MultiStore technology enables companies to better manage, consolidate, migrate, and replicate critical data with minimal effort and maximum return.

2. iSCSI Storage
   Data storage growth site.

7.2 Glossary

IP Space
IP space is a logical qualifier for IP addresses which allows a single storage system to properly handle possibly overlapping private IP addresses. Each interface belongs to only one IP space, but an IP space can have multiple interfaces.

8. Revision History

Date	Name	Description
03/31/2006	Kiran Sreenivasamurthy	Creation

© 2006 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, DataFabric, Data ONTAP, MultiStore, NearStore, SnapDrive, SnapMirror, SnapMover, and SnapVault are registered trademarks and Network Appliance, FlexVol, NOW, Snapshot, and vFiler are trademarks of Network Appliance, Inc. in the U.S. and other countries. Windows is a registered trademark of Microsoft Corporation.