IBM

N series Data Motion

- Data Motion in virtualized environments
- Data Motion in physical environments
- Dynamic Data Center enablement

Alex Osuna
Larry Touchette
Roger Weeks
Preetom Goswami

Redbooks

ibm.com/redbooks

International Technical Support Organization

N series Data Motion

August 2010

SG24-7900-00

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (August 2010)

This edition applies to Version 1 of Data Motion.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM®	Redbooks (logo)  ®	Tivoli®
Redbooks®	System Storage®	

The following terms are trademarks of other companies:

Vfiler, Snapshot, RAID-DP, FlexCache, SyncMirror, SnapVault, SnapMirror, SnapManager, SnapLock, SnapDrive, MultiStore, FlexVol, FlexClone, FilerView, Data ONTAP, NetApp, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

As customers and service providers begin to consolidate more and more applications and workloads onto shared storage infrastructures, it becomes increasingly difficult to coordinate outages for planned downtime for things, such as hardware refreshes. This difficulty is because many users, groups, or customers might be using the shared storage infrastructure at the same time. Users expect these infrastructures to be available 24/7, so it is imperative that service outages that are required for storage life-cycle management, cost/service-level optimization, and any other planned downtime do not disturb the availability of the *always-on* infrastructure.

In this IBM® Redbooks® publication, we introduce you to the business value of Data Motion and its features. Storage and system administrators, data center managers, and IT as a service (ITaaS) providers will benefit from reading this book.

In this publication, we assume that you have a basic knowledge of N series storage systems and Data ONTAP®, SnapMirror®, MultiStore®, and Provisioning Manager. A complete understanding of all Provisioning Manager features is not necessary. Also, basic knowledge of host and vFiler management from Provisioning Manager is sufficient.

Data Motion significantly improves the availability of shared storage infrastructure by avoiding the service outages that are associated with planned activities, such as storage life-cycle management and cost/service-level optimization, thus helping you to enable an *always-on* IT environment.

The business values of Data Motion are:

- No planned downtime for:
  - Storage capacity expansion
  - Scheduled maintenance outages
  - Technology refresh
- Improved SLA flexibility:
  - On-demand load balancing
  - Adjustable storage tiers
- Application transparency:
  - No performance impact
  - Transaction integrity

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Alex Osuna is a Project Leader at the IBM International Technical Support Organization, Tucson Arizona Center.

Larry Touchette is a Technical Marketing Engineer for NetApp®.

Roger Weeks is a Technical Marketing Engineer for NetApp.

Preetom Goswami is a Technical Marketing Engineer for NetApp.

Thanks to Eva Ho, IBM N series Product Engineer (PE), for her contributions to this project.

Thanks to the following people from NetApp for their contributions to this project:

Hamish McGovern, Paul Feresten, Bill L Heffelfinger,

Brian Gracely, Uday Shet, Jyoti Kurmi,

Nagender Somavarapu, Devender Mudhiganti, Srinivas Hemige,

Kaushik Sarkar, Suresh Kumar, Sridhara Gangoor,

Sitakanta Chaudhury, and Ratheesh Ramachandran

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author - all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

Introduction to Data Motion

This chapter provides an overview and describes the best practices for Data Motion.

The topics that we discuss in this chapter are:

- What is Data Motion?
- The components of Data Motion
- When to use Data Motion
- The software and hardware requirements of Data Motion
- Configurations supported by Data Motion
- How Data Motion works
- How to use Data Motion
- Limitations of Data Motion

1.1 Overview of N series Data Motion

This section describes the main benefits of IBM System Storage® N series Data Motion.

1.1.1 Business challenge

As customers and service providers begin to consolidate more and more applications and workloads onto shared storage infrastructures, it becomes increasingly difficult to coordinate outages for planned downtime for things, such as hardware refreshes, because many users, groups, or customers use the shared storage infrastructure simultaneously. Users expect these infrastructures to be available 24/7, so it is imperative that service outages that are required for storage life-cycle management, cost/service-level optimization, and any other planned downtime does not disturb the availability of the *always-on* infrastructure.

1.1.2 What is Data Motion

Data Motion is one of the foundation pieces of always-on data mobility, providing nondisruptive migration of MultiStore vFiler units for N series and MetroCluster. It is the data mobility component of cloud solutions.

Data Motion can move data from one Data ONTAP system to another while applications continue to run uninterrupted, as shown in Figure 1-1 on page 3. Data Motion also provides data mobility for multi-tenant environments because data is aggregated into secure virtual storage containers (vFiler units), which can correspond to a single customer, department, or application.

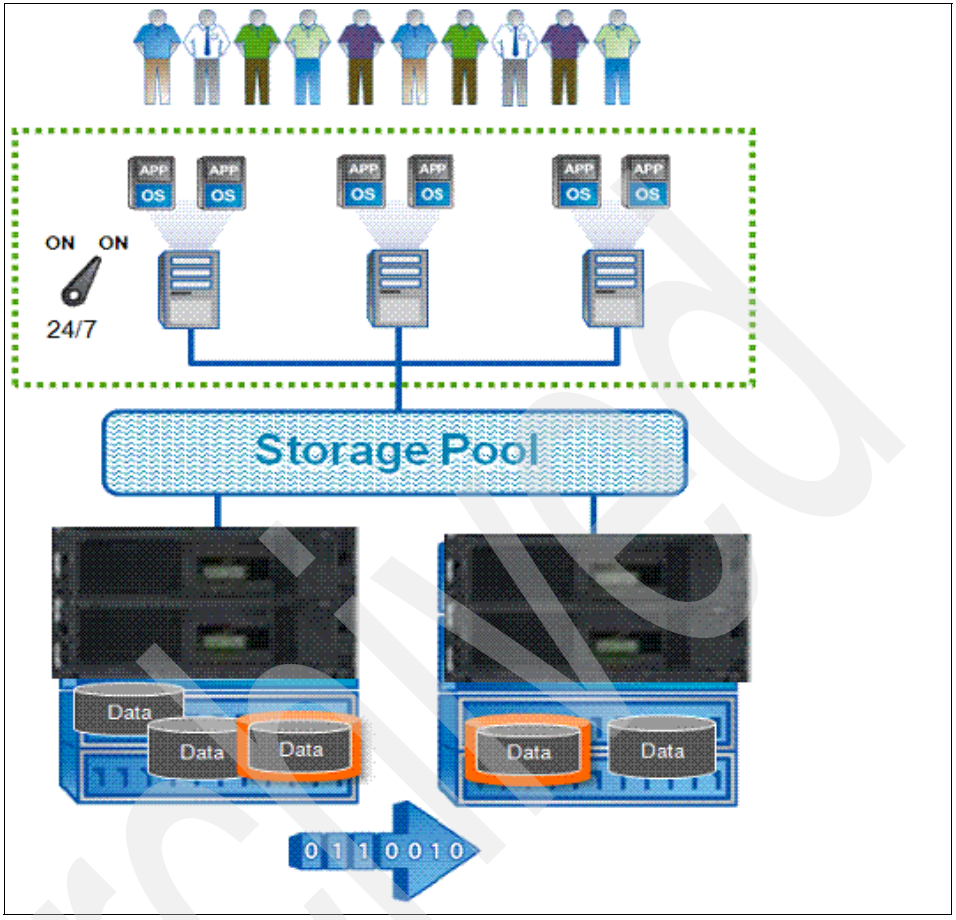

Figure 1-1 Data Migration

1.1.3 Business value of Data Motion

Data Motion significantly improves the availability of shared storage infrastructure by avoiding the service outages that are associated with planned activities, such as storage life-cycle management and cost/service-level optimization, thus helping customers to enable an always-on IT environment.

The business values of Data Motion are:

- No planned downtime for:
  - Storage capacity expansion
  - Scheduled maintenance outages
  - Technology refresh

- Improved SLA flexibility:
  - On-demand load balancing
  - Adjustable storage tiers
- Application transparency:
  - No performance impact
  - Transaction integrity

1.2 Data Motion use cases

In this section, we cover the environments in which Data Motion is most beneficial.

1.2.1 Data Motion in virtualized environments

Virtualization is driving large-scale consolidation of many applications in the same shared infrastructure, resulting in significant costs savings through greater utilization, reduced infrastructure, minimized floor space, and smaller power needs. It is now much more difficult to coordinate outages across many separate application teams or business units—and this is where Data Motion steps in. The combination of Data Motion with other mobility solutions that are generally available with virtual server infrastructure, such as VMware vSphere can provide nondisruptive data mobility at every layer of the virtual infrastructure.

Figure 1-2 on page 5 shows virtual machine migration.

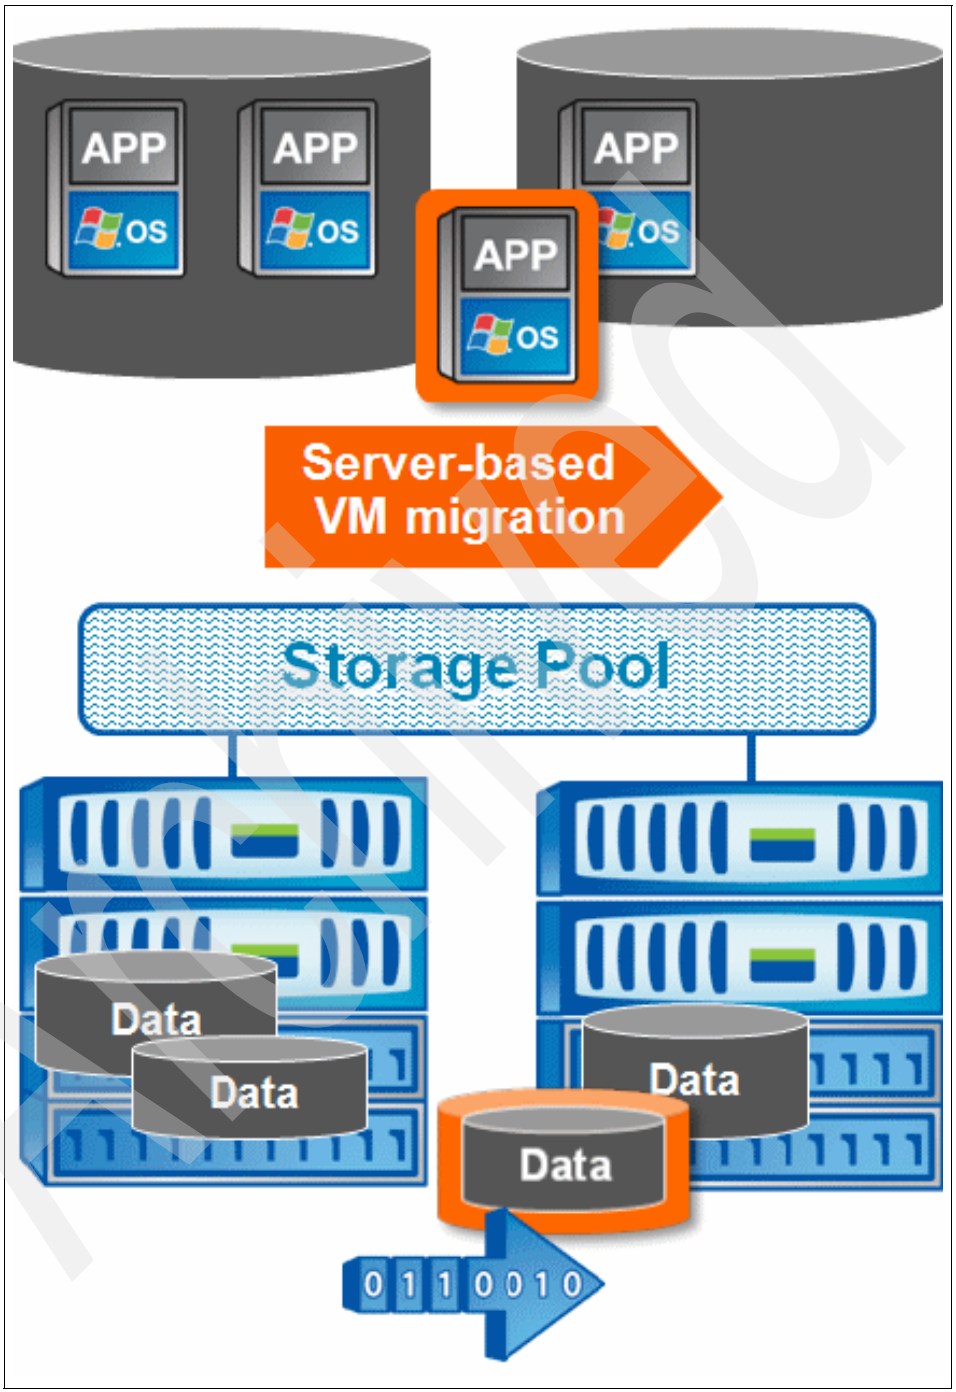

Figure 1-2 Virtual machine migration

Virtual server software, such as VMware vSphere, supports live migration of virtual machines across multiple server nodes in the cluster (VMotion), as shown in Figure 1-2 on page 5, and also provides the facility to move the virtual disks that are associated with the virtual machines online from one storage repository (or data store) to another (Storage VMotion). Data Motion complements migration technologies that are available from server virtualization vendors. It is important to recognize that the online migration technologies that Data Motion provides are designed to address the need for data mobility at a separate layer of the infrastructure and with a different effect on the data center environment.

Virtual server administrators use virtual-server-based migration to move individual virtual machines between data stores for purposes, such as capacity balancing between individual data stores. These migrations move one virtual machine at a time and consume host server resources, such as CPU cycles, memory, and network. Data Motion in a virtual server environment can be used to move entire data stores between storage arrays non-disruptively, for the reasons that we will discuss in 2.3, “Data Motion workflow and process” on page 16. Enterprise storage administration teams typically perform these operations.

Table 1-1 provides the online migration techniques that can be used in a virtual infrastructure.

Table 1-1 Online migration techniques

Migration scenario	Tool to be used	Use case	Used by	Typical occurrence	Target customer
Virtual machine management (migration of individual VMs)	VMware Storage VMotion	Movement of virtual disks that are associated with a VM	Virtual Admin	Daily or Weekly	Any customer using a virtual infrastructure
Storage infrastructure management (migration of entire datastore)	Data Motion	Movement of an entire datastore, customer, department, or application	Storage admin	Monthly	Enterprise, service provider, ITaaS, PaaS

1.2.2 Data Motion in physical environments

This section covers how Data Motion reacts in physical environments.

Database and messaging

Databases and messaging systems are the heartbeat of most companies with nearly every staff member using them on a constant basis. As businesses become more global, outages on these systems, even after hours, not only affect many individuals, but restarting the applications can take much longer than the outage itself. The ability to move the application data from one array to another without shutting down the application is extremely compelling.

Web applications

Externally-facing web applications are in use all of the time, so routine maintenance can mean lost business with customers on the web. Web users expect immediate gratification, and if a site is down, they might go elsewhere. Data Motion removes the need for outages for life-cycle management, maintenance, and capacity and performance management of web application environments.

1.2.3 Enabling the dynamic data center

Data Motion can be considered as one of the enablers of the dynamic data center, where avoiding planned downtime in any part of the infrastructure is of paramount interest. In the following subsections, we describe use cases of Data Motion in this context.

Secure multi-tenancy

Data Motion is inherently enabled for multi-tenancy because it is based on MultiStore technology. This technology provides mobility at the level of individual customers, departments, or application data sets that must be isolated and secure.

Technology refresh

For simplified life-cycle management, new storage technology can be accomplished without the need for planned downtime. Installation of new N series storage systems and higher capacity disk drives or shelves can be done non disruptively.

For nondisruptive movement of data to a newer platform, the old platform must run at least Data ONTAP 7.3.3.

Dynamic capacity scaling

Where application data exceeds the storage capacity of an existing N series storage system or N series Gateway system, new storage systems can be

installed, and existing data can be transitioned to the new storage without any disruption to applications.

Service-level optimization

For flexible service-level optimization (SLA) management, Data Motion can be used to move data across N series storage systems for load balancing. Data can also be moved across storage tiers (slower disks to faster disks) to optimize performance.

2

Components of Data Motion

In this chapter, we discuss the Data Motion components. The topics that we discuss are:

- MultiStore
- Provisioning Manager
- SnapMirror
- Supported platforms
- Hardware and software requirements

2.1 Overview of N series Data Motion components

Data Motion is a solution that is enabled by the combination of features from three N series products:

- MultiStore

  Provides the capability to partition a physical N series storage system into virtual systems called vFiler units, which are the units of migration in Data Motion. MultiStore enables secure logical partitioning of network and storage resources in Data ONTAP into vFiler units.

- Provisioning Manager

  Provides the software interface for performing Data Motion migrations.

- SnapMirror

  Data is migrated by using SnapMirror. SnapMirror asynchronous mode is used for the initial baseline transfer phase of Data Motion. However, after the migration cutover process starts (that is, when the control of the vFiler unit actually migrates from one physical storage system to another), SnapMirror switches to a semi-synchronous mode of transfer.

- Flexible Volumes (FlexVols)

  FlexVols are file systems that hold user data that is accessible through one or more of the access protocols that Data ONTAP supports.

2.1.1 MultiStore

MultiStore software provides separate and completely private logical partitions on a single N series storage system as discrete administrative domains called vFiler units. These vFiler units have the effect of making a single physical storage controller appear to be many logical controllers. Each vFiler unit can be individually managed with different sets of administrative and policy characteristics, as shown in Figure 2-1 on page 11.

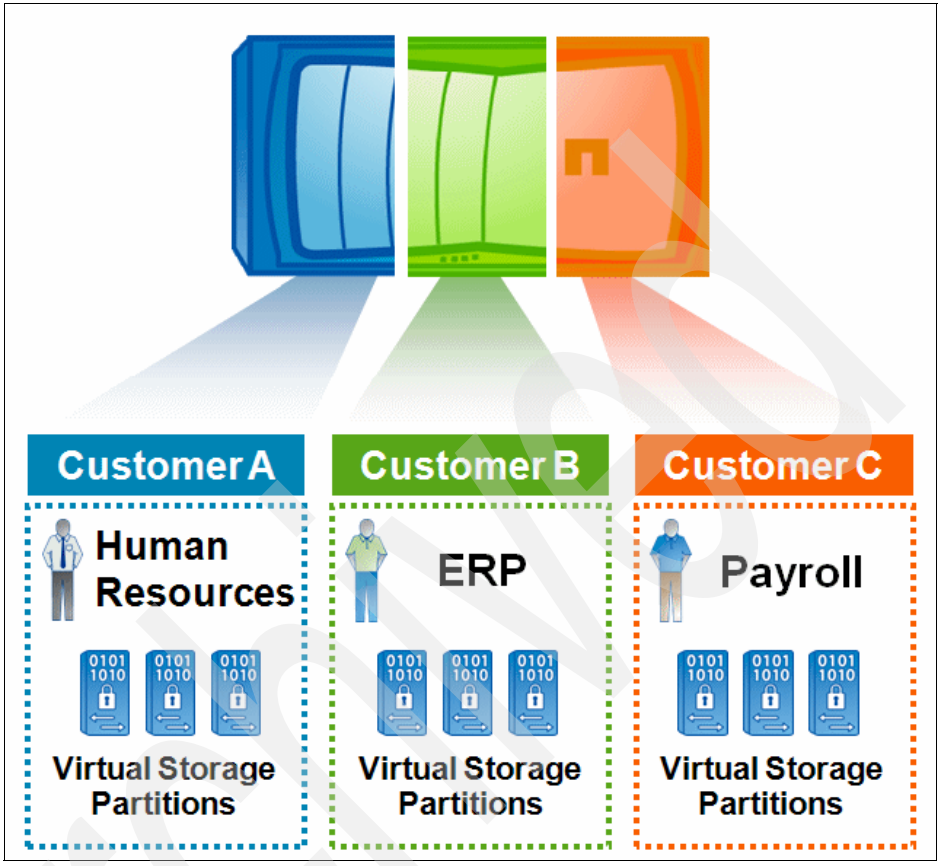

Figure 2-1 Multistore use example

Technology service and cloud providers can use MultiStore to enable multiple customers to share the same storage resources with no compromise in privacy or security. They can also delegate administrative control of the vFiler unit directly to the customer.

2.1.2 Provisioning Manager

Provisioning Manager is a feature in Operations Manager that is activated with a license key. Provisioning Manager enables storage administrators to use policy-based automation to create repeatable, automated provisioning processes. These processes are faster than manually provisioning storage, easier to maintain than scripts, and help to minimize the risk of data loss because of incorrect storage configuration.

Provisioning Manager is responsible for the end-to-end workflow of Data Motion processes.

2.1.3 SnapMirror: Asynchronous and synchronous

SnapMirror, shown in Figure 2-2, is a data replication feature that efficiently mirrors data between two N series storage systems. The source and destination controllers can live in the same data center or be geographically separated because both LAN and WAN links are supported. SnapMirror runs at scheduled intervals to keep the destination controller up-to-date, so the recovery point objective (RPO), which is the amount of data not yet backed up, is dramatically lower than with mature, periodic backup routines. A SnapMirror relationship can also be placed into synchronous or semi-synchronous mode, which effectively reduces the RPO to almost zero. A N series destination storage system is more than a simple backup or replica depot because it can also serve as the source system in a disaster recovery scenario, which means that the recovery time objective (RTO), which is the time from disaster to being back online, is also minimal.

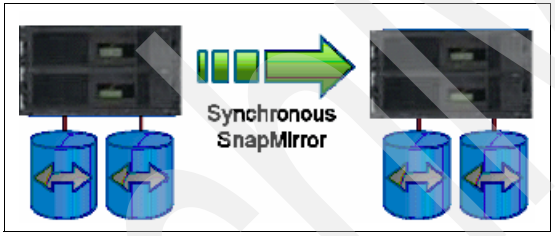

Figure 2-2 Synchronous SnapMirror

Data Motion transfers are initially set up in SnapMirror asynchronous mode; however, after the baseline transfer is complete and the cutover phase is initiated, the transfers are performed in SnapMirror semi-synchronous mode. During the cutover phase, every FlexVol® volume that a vFiler unit owns must be migrated in semi-sync mode simultaneously.

2.1.4 Flexible Volumes

FlexVols are file systems that hold user data that is accessible through one or more of the access protocols that Data ONTAP supports, which includes Network File System (NFS), Common Internet File System (CIFS), HTTP, FTP, FCP, and iSCSI. Because each flexible volume is a separate file system, you can create one or more Snapshots of the data in a volume so that multiple, space-efficient, point-in-time images of the data can be maintained for purposes, such as backup and error recovery.

FlexVol technology is a ground breaking technology that comes embedded with DATA ONTAP software. FlexVols are independent of the underlying physical storage. These are the logical entities that are sized, resized, managed, and moved independently of the underlying storage.

Volumes remain the primary unit of data management. Flexible volumes refer to logical entities, not (directly) to physical storage, and are transparent to the administer.

The physical storage that supports flexible volumes is first arranged in RAID groups (either RAID 4 or RAID-DP™, which is the default). Refer to Chapter 5, “Data Motion best practices” on page 49, for more information about this topic.

One or more RAID groups are then combined together into an *aggregate*, as illustrated in Figure 2-3.

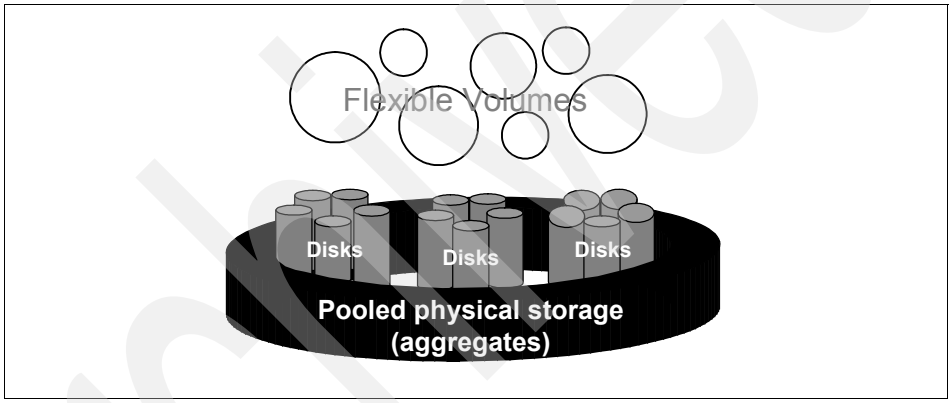

Figure 2-3 Flexible volumes

Each storage appliance can support multiple aggregates, with the maximum number dependent on the capacity of the storage appliance and the version of Data ONTAP.

Each volume depends on its containing aggregate for all of its physical storage (that is, for all storage in the aggregate's disks and RAID groups), as illustrated in Figure 2-4.

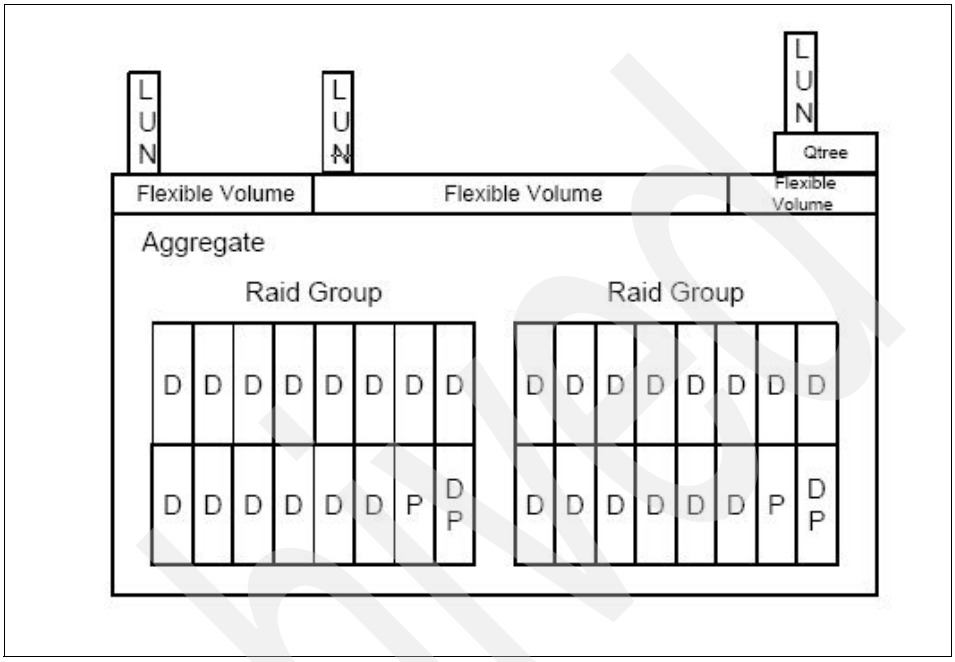

Figure 2-4 Flexible volumes: Storage

Because the FlexVol is managed separately from the aggregate, you can create small FlexVol volumes (20 MB or larger), and you can increase or decrease the size of FlexVol volumes in increments as small as 4 KB.

A FlexVol can share its containing aggregate with other FlexVol volumes. Thus, a single aggregate can be the shared source of all of the storage that is used by all of the FlexVol volumes that are contained by that aggregate. The aggregate manages the unused space so that unallocated space in one FlexVol does not impact the space that is used in another FlexVol within the same aggregate.

As illustrated in Figure 2-5 on page 15, an aggregate is defined as a pool of many disks from which space is allocated to volumes. (Volumes are shown in Figure 2-5 on page 15 as FlexVol and FlexClone® entities.)

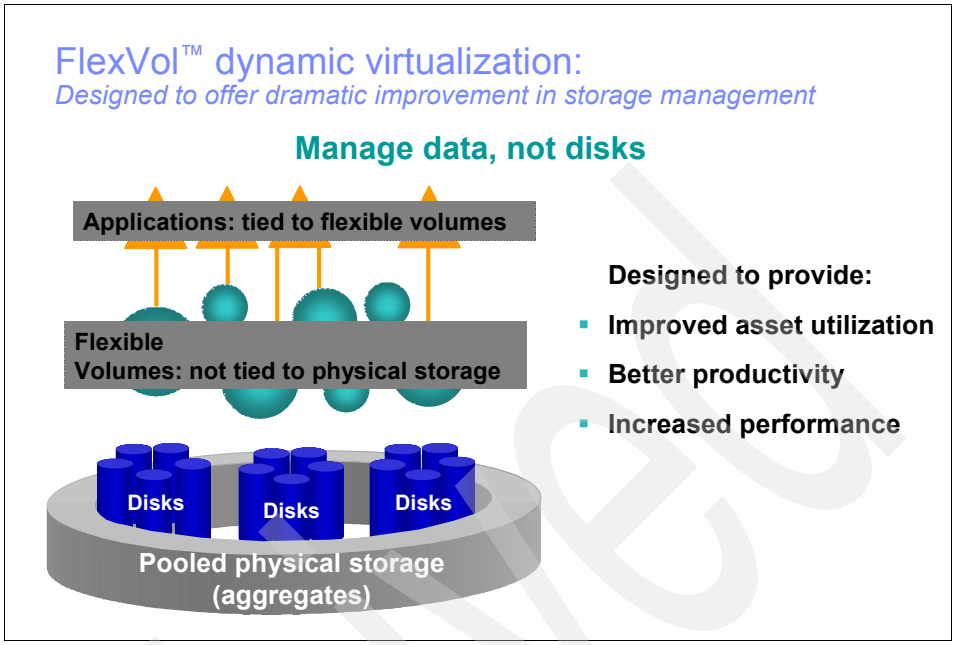

Figure 2-5 An aggregate: A pool of many disks from which space is allocated to volumes

From the administrator's point of view, volumes remain the primary unit of data management. But transparently to the administrator, flexible volumes now refer to logical entities, not (directly) to physical storage.

2.2 Supported platforms and requirements

This section lists the software and hardware requirements for the Data Motion components.

2.2.1 Software requirements

This section covers the software requirements:

- Data ONTAP 7.3.3 or higher

  Contains vFiler migration enhancements and APIs that Data Motion requires.

- Supported Protocols:
  - NFS or iSCSI: No FCP or FCoE support at this time
  - CIFS requires clients to reconnect after the migration

- Data ONTAP Licenses:
  - NFS/iSCSI
  - MultiStore
  - SnapMirror, SnapMirror Sync
- Provisioning Manager 4.0
- Operations Manager 4.0

2.2.2 Hardware requirements

This section covers hardware requirements:

- N3000[1], N5000, N6000, and N7000 series
- N series Gateway systems—N7000, N6000
- MetroCluster platforms

2.3 Data Motion workflow and process

Provisioning Manager provides the user interface for the Data Motion process and also controls the entire workflow. For an end-to-end Data Motion scenario, see Appendix A, “The Data Motion process” on page 57.

Although Provisioning Manager is required to perform a Data Motion migration, it is not required to use Provisioning Manager to create or provision the vFiler units that are migrated with Data Motion. Operations Manager and Provisioning Manager automatically detect any pre-existing vFiler units in the environment or vFiler units that were created with other tools or processes.

2.3.1 Advantages of using Provisioning Manager for Data Motion workflow control

This section covers the advantages of Provisioning Manager:

- Enforcement: Provisioning Manager automatically enforces Data Motion requirements. It aids successful migrations and helps prevent application outages.
- Automation and ease of use: Manual completion of the Data Motion steps, such as establishing and starting the SnapMirror transfers, post migration cleanup, and so on, is not necessary. Example scenarios that might disrupt a

[1] Excluding N3700

manual online migration process, which can be avoided by using Provisioning Manager are:

- Too much load on the systems that are involved
- Configuration changes that are made between the SnapMirror initialize and the cutover (for example, new networks or volumes added)
- Other SnapMirror relationships performing transfers and using transfer queues

2.3.2 Phases of Data Motion

A Data Motion workflow can be divided into three high-level phases:

- Initial baseline transfer
- Cutover
- Cleanup

Note: After the cutover phase, it is possible to quickly rollback the vFiler unit to the original source if the cleanup phase is not yet executed. This process provides rapid online migration of the vFiler unit from the destination storage system to the source storage system without going through the initial baseline transfer phase again.

Initial baseline transfer phase

In this *initial baseline transfer phase*:

1. The Data Motion process is initiated from the Provisioning Manager interface.
2. Provisioning Manager performs the following steps:
   a. Verifies that systems meet Data Motion requirements.
   b. Creates destination VLAN interfaces and IP spaces.

      (VLANs must be enabled on destination switch ports by the administrator.)
   c. Creates destination FlexVol volumes.
   d. Initiates SnapMirror baseline transfers.
   e. SnapMirror schedule is set to asynchronously update every three minutes.
   f. When baseline completes, migration status is set to Cutover Required, as shown in Figure 2-6 on page 18.

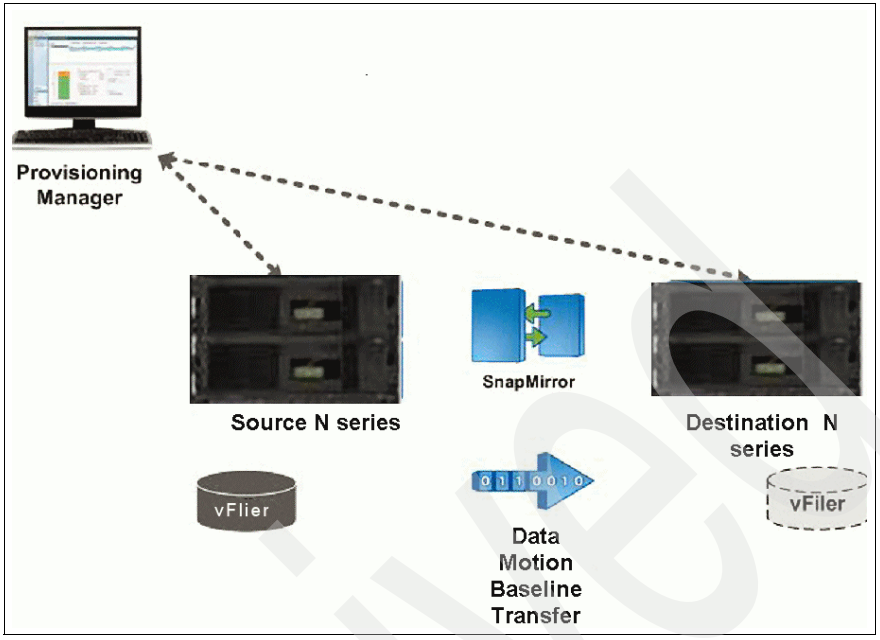

Figure 2-6 Baseline transfer

Note: While starting the initial baseline transfer process from the Provisioning Manager, there is an option to cutover automatically after data transfer. This option enables one step migration of the vFiler unit where the N series storage system automatically performs the cutover operation after the initial baseline transfer is completed. By default, this option is not selected. So under a default scenario, you must manually initiate the cutover operation at any time after the data is transferred.

After the initial baseline transfer is complete and the migration status is changed to `Started, cutover required`, you have the option to perform an Update before proceeding to the next state (Cutover). The migration Update feature must be used if there is a long gap between the time when initial baseline transfer was completed and the time when the cutover phase is initiated.

Cutover phase

During the *cutover phase*:

1. The cutover is either initiated manually from Provisioning Manager, or it is started automatically (if the automatic cutover option is selected at the start of the Data Motion process).

2. Provisioning Manager performs the following steps:
   a. Verifies that systems meet Data Motion requirements.
   b. Checks the vFiler unit for configuration changes since migration start.
   c. Converts the SnapMirror relationships from asynchronous mode to semi-synchronous mode.
   d. Starts transparent vFiler migration.
   e. If CIFS is enabled on the source vFiler unit, it is stopped, then restarted after cutover.

   Data Motion fences admin operations on the vFiler volumes before cutover. Some typical admin operations that are fenced are: Snapshot™ creation, file/LUN FlexClone creation, SnapMirror or SnapVault® updates on the vFiler volumes, and so on, which is required for smooth cutover operation. This fencing continues until cutover completes.
3. The cutover process is handled by Data ONTAP, as shown in Figure 2-7 on page 20:
   – During the cutover, the migration must complete within the 120-second window. The two Data ONTAP 7.3.3 storage systems negotiate with each other to perform the migration. Provisioning Manager is not involved in this part of the process; therefore, its failure during this time does not cause a failure of the Data Motion migration or an application outage.
   – Data ONTAP maintains timers that watch the vFiler migration process.
   – If the cutover fails or exceeds the 120-second window, the process is aborted and the original vFiler unit is restarted within the same 120-second window.

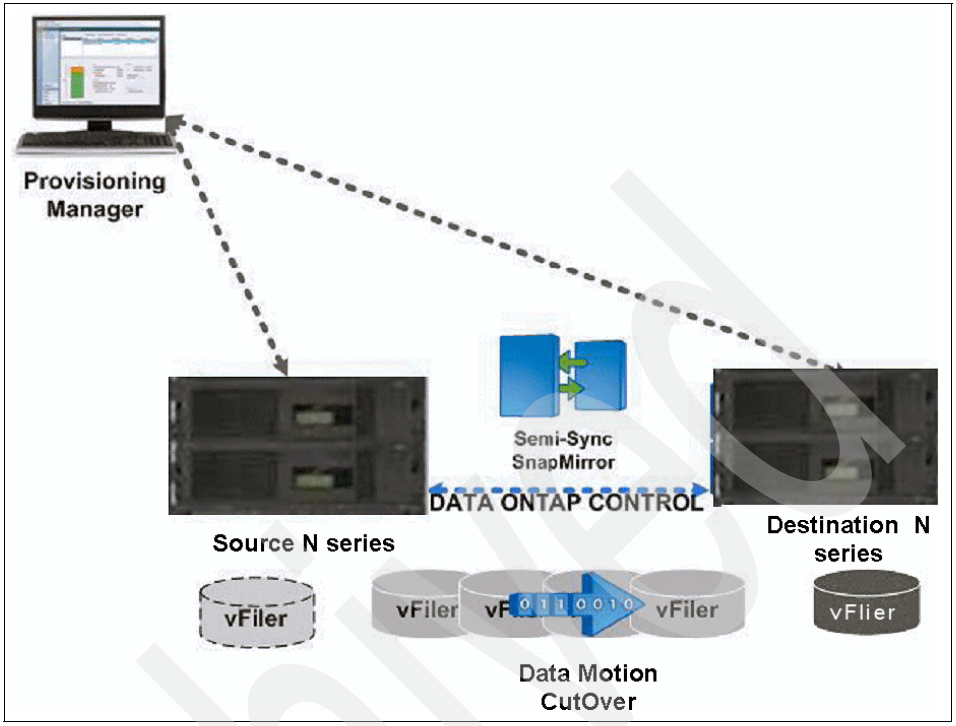

Figure 2-7 Cutover

Client experience during online migration

The client experience during online migration includes:

- iSCSI clients experience up to 120-second pause in I/O.
- NFS clients experience up to 120-second pause in I/O.
- Clients reestablish NFS locks as they do during a high-availability takeover.
- CIFS client sessions are terminated after cutover and the clients must reconnect.
- FCP clients are not supported by Data Motion at this time.

After successful completion of Data Motion cutover

After successful completion of the Data Motion cutover phase, Provisioning Manager automates the following post-migration tasks:

1. Migrates any Protection Manager relationships.
2. Modifies Provisioning Manager data set memberships.
3. Migrates Operations Manager history.

4. NDMP backups from vFiler unit (not direct to tape) continue as normal.
5. The original source volumes are maintained in an offline state to support the optional rollback feature.

Additionally, NDMP backups from vFiler0 that correspond to the vFiler volumes must be reconfigured.

Data Motion cutover failure

If the Data Motion cutover fails:

1. Data ONTAP makes sure that the original vFiler unit is brought back online.
2. Provisioning Manager performs these post-failure steps:
   a. SnapMirror relationships are maintained for cutover retry.
   b. SnapMirror relationships are set back to async mode.
   c. User can later retry the cutover from Provisioning Manager.
   d. If migration is canceled after failure, Provisioning Manager cleans up all new SnapMirror relationships, VLAN interfaces, IP spaces, FlexVol volumes, and so on.

Cleanup phase

The *cleanup phase* is executed only when it is initiated from Provisioning Manager, as shown in Figure 2-8 on page 22.

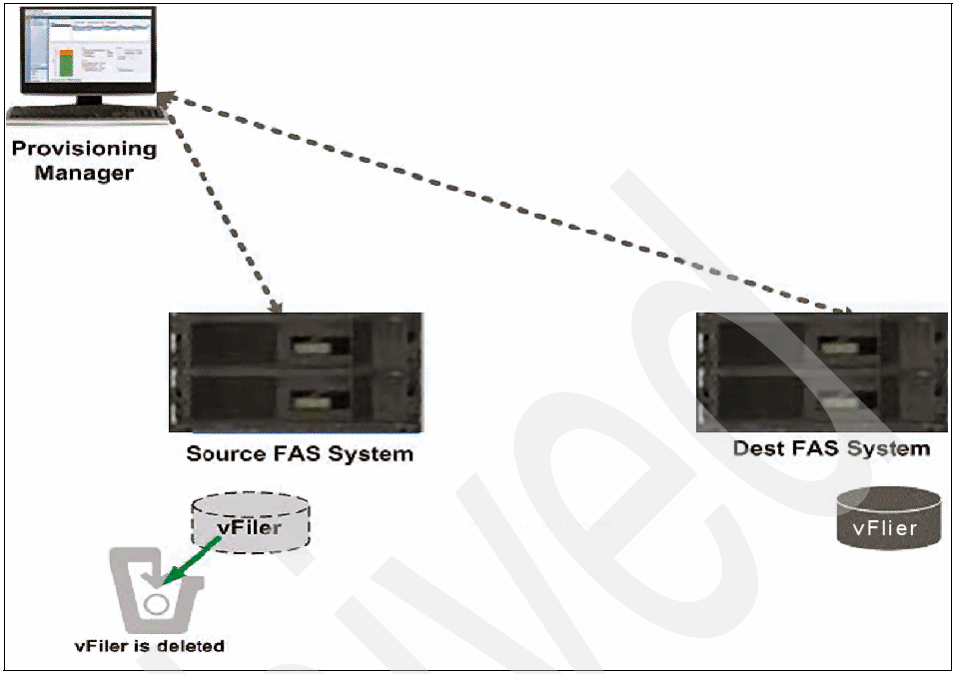

Figure 2-8 Cleanup Phase initiated from Provisioning Manager

During the cleanup phase, Provisioning Manager automatically removes all source FlexVols.

Rollback feature (optional)

The Rollback feature of Data Motion provides the ability to quickly migrate the vFiler unit back to the source controller after a successful cutover. This feature can be used only when the administrator has not executed the cleanup phase.

This feature might be useful if you face any unforeseen issues after cutover (such as, performance degradation) and want to revert to the original configuration quickly without performing the initial baseline transfer phase again.

Note: As mentioned in the Supported Configuration section, Data Motion is supported only between similar system models (systems having the same size NVRAM) or from a low-end system model to a high-end model (systems with smaller size NVRAM to a system with larger size NVRAM). This rule applies to the Rollback feature too, which means that a vFiler unit can be rolled back after cutover only if the migration is between two similar system models (or systems with same size NVRAM). In other words, for a Data Motion migration from a low-end system to a high-end system, the rollback feature is not available.

Data Motion is supported only between disks of the same RPM speed or from slower to faster disks. This rule also applies to the rollback feature, which means that a vFiler unit can be rolled back after cutover only if the migration is between disks of the same RPM speed.

Other features of Data Motion workflow

Additional features of the Data Motion workflow are:

- All backup relationships are migrated when used with Protection Manager.
- AutoSupport reports are generated for any Data Motion failures that have an impact on application availability, such as cutover exceeding the 120 sec window or the source vFiler unit not coming up within the 120 sec window after an aborted cutover process.

Summary of Data Motion steps

Figure 2-9 on page 24 illustrates the steps from the initial baseline transfer phase to the cleanup phase.

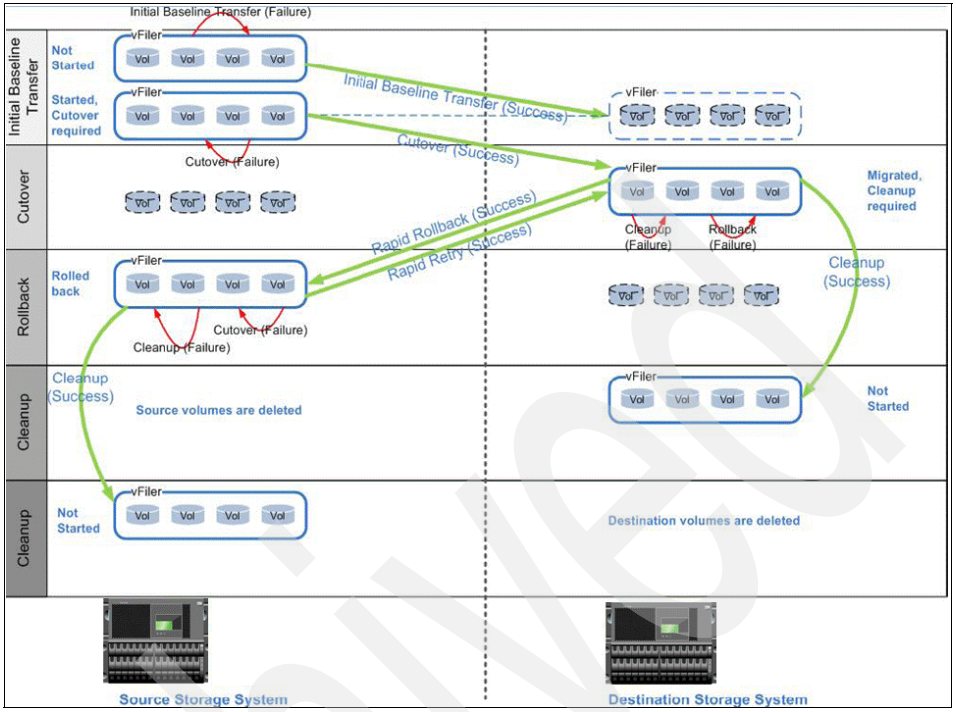

Figure 2-9 Data Motion steps

3

Supported configurations and requirements

In this chapter, we cover supported configurations and the requirements for Data Motion. The topics that we discuss in this chapter are:

- Supported migration scenarios
- Supported vFiler configurations

3.1 Supported migration scenarios

Table 3-1 shows an example of supported migration scenarios.

Table 3-1 Supported MIgration scenarios

Migration Scenario	Description	Supported?
	Between controller heads of different high-availability pairs	Yes
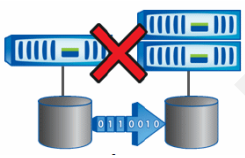	Between a single controller head and a controller head of a high-availability pair	No
	Between the controller head of a high-availability pair and a single controller head	No
	Between two single controller heads	No
	Within a controller (between aggregates)	No

Migration Scenario	Description	Supported?
	Between controller heads of the same high-availability pair	No

3.1.1 Non Supported Data Motion Migration configurations

In this section, we identify Data Motion Migration unsupported configurations:

- Data Motion supports migration of vFiler units from a N series node that is part of a N series high-availability pair to another node that is part of a different N series high-availability pair. Both source and destination high-availability clusters must be up and running. If source or destination nodes are in takeover/failback mode, Data Motion is not supported and execution of this scenario is prevented by Provisioning Manager.
- Data Motion vFiler migration between the two nodes of a N series high-availability pair is not supported and execution of this scenario is prevented by Provisioning Manager.
- Data Motion migration from a single node N series to another single node N series (which are not part of the any high-availability pair) is not supported. Although Provisioning Manager does not prevent the execution of this scenario, this migration scenario is not recommended.
- Data Motion migration is not supported from a single node N series to another node that is part of a N series high-availability controller pair and vice versa. Although Provisioning Manager does not prevent the execution of these scenarios, these migration scenarios are not recommended.

3.1.2 Configuration support

Data Motion is supported only between similar N series Models (that is between systems having the same amount of NVRAM) or from a low-end system model to a high-end system model (that is from a system with smaller NVRAM to a system with larger NVRAM), as shown in Table 3-2 on page 28.

Table 3-2 Data Motion scenarios

Migration scenario	Description	Supported?
	Between two high-availability pairs having the same NVRAM	Yes
	From a high-availability controller having smaller NVRAM to a high-availability controller having larger NVRAM	Yes
	From a high-availability controller having smaller NVRAM to a high-availability controller having smaller NVRAM	No

3.1.3 Drive speed requirements

Data Motion is supported only between same speed drives or from slower to faster speed drives, as shown in Table 3-3 on page 29.

Table 3-3 Drive speed requirements

Migration scenario	Description	Supported?
	Between same-speed drives	Yes
	From slower to faster drives	Yes
	From faster to slower drives	No

3.1.4 Vfiler support

Data Motion is supported only for Vfiler™ units that own FlexVols as storage containers, as shown in Table 3-4 on page 30.

Table 3-4 Vfiler support

Migration Scenario	Description	Supported?
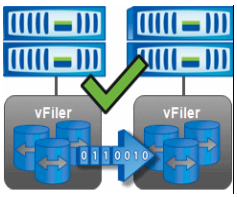	vFiler units owning FlexVol volumes where an owned FlexVol volume might contain qtrees	Yes
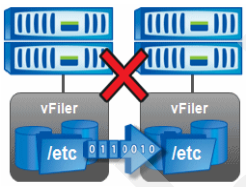	► vFiler units owning qtrees whose owning FlexVols reside outside the vFiler unit ► vFiler units owning traditional volumes	No

3.2 Supported vFiler configurations

Data Motion performs migrations by migrating entire vFiler units, so all FlexVols that are owned by the vFiler unit must be migrated together in one operation. Table 3-5 on page 31 lists the maximum number of FlexVols per vFiler unit, supported by Data Motion corresponding to various storage platforms.

Data Motion restrictions: These restrictions were chosen for the first release of Data Motion based on tests that were performed to verify successful migrations under certain load conditions. As more conditions are qualified, these limitations can be relaxed.

Table 3-5 FlexVols per vFiler

Platform	FlexVols per vFiler unit
N 3000[a]	4 FlexVols
N 5000[b]	8 FlexVols
N 6000	8 FlexVols
N 7000	20 FlexVols

a. N 3700 not supported
b. Data Motion supports 4 FlexVols for N 5500 and N5200 platform.

3.2.1 FlexVol requirements for Data Motion

In this section, we discuss the aspects of FlexVol that are required for Data Motion.

Size of FlexVol volumes

When working with vFilers, Data Motion has the following FlexVol requirements:

- Data Motion requires FlexVols that a vFiler unit owns to be at least 10 GB in size. This is a requirement of synchronous SnapMirror.
- The vFiler unit's root volume is typically less than 10 GB. In that case, Data Motion automatically increases the size of the FlexVol to 10 GB, performs the migration, and then resizes the volume to its original size.

Size of Aggregates

To achieve the maximum number of FlexVols per vFiler unit, it is recommended that aggregates contain at least 12 disks.

SnapLock Not Supported

SnapLock® volumes cannot be migrated with Data Motion.

Support for CIFS

CIFS shares within FlexVols that are owned by the vFiler unit are migrated as part of Data Motion. However the CIFS sessions are terminated and clients must reconnect. Unlike NFS and iSCSI, which are stateless protocols, CIFS is a session-oriented protocol, which can result in adverse effects on clients and applications during migrations. Instruct users to end their sessions before the cutover.

Support for deduplication and File/LUN FlexClone

Deduplication is an N series storage efficiency offering that provides block-level deduplication within the entire flexible volume on N series storage systems. Beginning with Data ONTAP 7.3, N series gateways also support deduplication. N series gateways are designed to be used as a gateway system that sits in front of third-party storage, allowing N series storage efficiency and other features to be used on third-party storage. Figure 3-1 shows how N series deduplication works at the highest level.

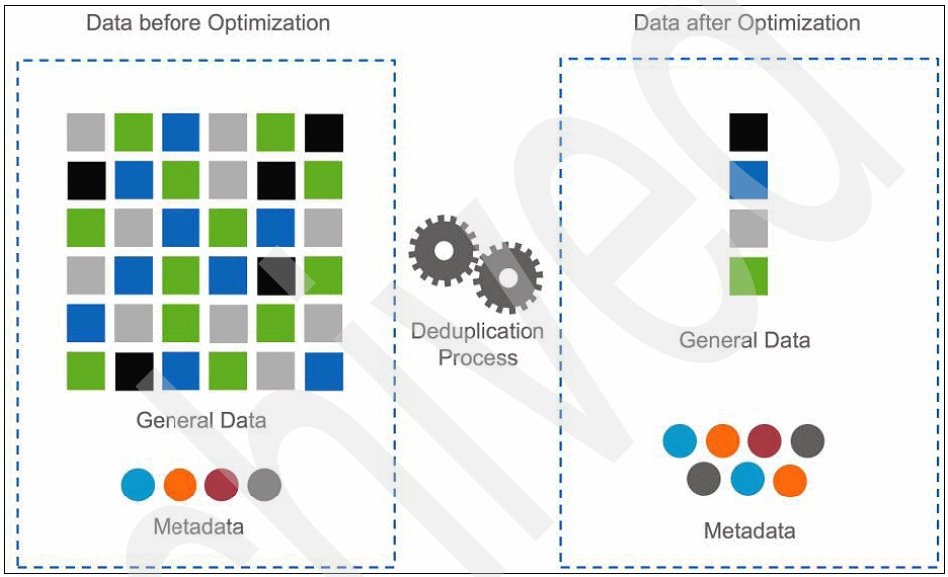

Figure 3-1 Deduplication

When Data Motion is used in conjunction with deduplication you should consider the following information:

- Data Motion is supported with dedupe or FlexClones at the Files or LUN level. However, the combination of having to manage resources that are associated with nondisruptive migration and metadata for dedupe or FlexClones can result in a small possibility of degradation to the performance of client applications, which can happen only if Data Motion is performed under high load (greater than 60%). Therefore, Data Motion must not be attempted under such conditions.
- It is recommended that Data Motion be performed in off peak or periods of lower load to ensure the fastest migration times and hence minimal impact. In this initial release, it is recommended that customers actively monitor the system performance during Data Motion cutover operation for systems that have dedupe or FlexClones at the Files or LUN level (on either the source or

destination system). If excessive load is encountered during the migration, Data Motion can be aborted by the storage administrator, while maintaining access to data.

- After a successful Data Motion cutover, the deduplicated volume remains deduplicated on the destination array, but before another dedupe process can run on the volume, the dedupe fingerprint database must be rebuilt. The Data Motion software automatically starts rebuilding the fingerprint database after a successful migration.

Support for Volume FlexClone and LUN clone

FlexClones can be members of vFiler units that are migrated with Data Motion. FlexClones are expanded at the destination, meaning that they are full-size independent volumes at the destination. Table 3-6 describes three FlexClone scenarios and their supportability with Data Motion.

Data Motion supports LUN Clones that are contained inside the volumes that are members of the vFiler unit that is being migrated.

Table 3-6 LUN Clone Support

Parent Volume	FlexClone Volume	Data Motion supported?
Inside the vFiler unit being migrated	Inside the vFiler unit being migrated	Yes FlexClone volumes are expanded at the destination
Outside the vFiler unit being migrated (that is, inside another vFiler unit or in vFiler0)	Inside the vFiler unit being migrated	No
Inside the vFiler unit being migrated	Outside the vFiler unit being migrated (that is, inside another vFiler or vFiler0)	No

3.2.2 Using Data Motion with other products

Table 3-7 on page 34 shows the interoperability of Data Motion with other products.

Table 3-7 Other product support

Products	Supported for Data Motion	Comments
Operations Manager	Yes	Provisioning Manager feature in Operations Manager controls the work flow of Data Motion.
SnapDrive® for UNIX/Linux® v4.x	Yes	
SnapDrive for Windows® v6.1	Yes	
Flexible Volumes	Yes	
Fractional reservations	Yes	
volume FlexClone	Partial	See 3.2.1, “FlexVol requirements for Data Motion” on page 31.
LUN Cloning	Yes	
Snapshot	Yes	
Fabric MetroCluster	Yes	
Qtree SnapMirror	Yes	Operation is fenced during Data Motion Cutover.
Volume SnapMirror Async	Yes	Operation is fenced during Data Motion Cutover.
SyncMirror®	Yes	
SnapManager® for Oracle 3.0.3	Yes	
SnapManager for SAP 3.0.3	Yes	
SnapManager for Virtual Infrastructure (SMVI) 2.0	Yes	
SnapManager for Hyper-V	Yes	
Deduplication	Partial	See “Support for deduplication and File/LUN FlexClone” on page 32.

Products	Supported for Data Motion	Comments
File/LUN FlexClone volumes	Partial	See “Support for deduplication and File/LUN FlexClone” on page 32.
NDMP/Dump	Yes	Operation is fenced during Data Motion Cutover.
PAM 1	Yes	The working set that is cached in the PAM card of source storage system for the vFiler volumes is not transferred to the destination storage system during Data Motion. So after Data Motion completes, the working set must be populated again on the destination storage system (for example, PAM cache must be *warmed* back up) and throughput can degrade during this time.
PAM II	Yes	
Volume SnapMirror Sync	Partial	Only in default vFiler interfaces. Also, operation is fenced during Data Motion cutover.
SnapVault	Partial	A SnapVault source can be a vFiler unit, but the destination must not.
FlexCache™	No	Only FlexCache source volumes can be migrated with Data Motion.
SnapLock Compliance	No	SnapLock volumes are not supported for Data Motion.
SnapLock Enterprise	No	SnapLock volumes are not supported for Data Motion.
SnapManager 5.0 for Microsoft Exchange	Yes	
SnapManager 5.0R1 for Microsoft SQL Server	Yes	
SnapManager for Microsoft Office SharePoint Server	Yes	
Thin Provisioning	Yes	

4

Network requirements

This chapter provides the network requirements for Data Motion. The topics that we discuss in this chapter are:

- IP requirements for Data Motion
- Other networking considerations

4.1 IP requirements for Data Motion

With the vFiler unit being the unit of migration in Data Motion, all MultiStore specific networking requirements must be met.

There are several ways to assign network addresses to vFiler units. A vFiler unit has interfaces that are assigned to it as part of the provisioning process. These interfaces can be any of the following items: physical interfaces, IP aliases, VLANs, or virtual interfaces (VIFs)—either using the default or a non-default IP space. It is not recommended to use physical interfaces because the system quickly runs out of available interfaces. Figure 4-1 illustrates the various network configuration options for vFiler. Options 4 and 8, which are emphasized with blue ovals, are the recommended vFiler network configuration options for Data Motion.

Options →	1	2	3	4	5	6	7	8
Physical Interface	•		•		•		•	
VIF		•		•		•		•
IP Alias	•	•	•	•	•	•	•	•
VLAN			•	•			•	•
IP Space	Default	Default	Default	Default	Non Default	Non Default	Non Default	Non Default

Figure 4-1 Network configuration options

4.1.1 Using IP aliases

For assigning IP addresses to the vFiler units, multiple IP aliases are created on a network interface of the physical N series storage system. Each vFiler unit is assigned an IP alias. All clients of the vFiler unit must make the request on this IP alias. This method is recommended for situations where vFiler units do not need physical or logical network separation.

4.1.2 Using VLANs and VIFs

Both types of virtual interfaces, VLANs and VIFs, can be assigned to a vFiler unit. Although you can assign an IP address directly to a VIF, this is not the method used in most MultiStore deployments. A VIF is created, and a VLAN is attached to that VIF. The VLAN is then assigned to an IP space, and a vFiler unit is created in that IP space.

4.1.3 Using the IPSpace feature

MultiStore introduces the concept of an IPSpace for supporting secure multi-tenancy. IPSpaces are distinct IP address spaces with private routing tables in which vFiler units reside. Each vFiler unit can reside in only one IPSpace; however one IPSpace can have multiple vFiler units. An IPSpace defines a distinct IP address space in which vFiler units can participate. IP addresses that are defined for an IPspace are applicable only within that IPSpace. No cross-IPSpace traffic is routed.

The requirements for the IPSpace feature are:

- Each network interface (physical, VIF, or VLAN interface) of the N series system can belong only to one IPSpace; however one IPSpace can have multiple network interfaces.
- IPSpaces are created when vFiler units must have their own secure storage, administration, and routing leading to secure multi-tenancy.
- IP addresses that are defined for an IPSpace are accessible only within the network of that IPSpace.
- A distinct routing table is maintained for each IPSpace. No cross-IPSpace traffic is routed.
- IPSpace has a unique loopback interface assigned to it. The loopback traffic of each IPSpace is completely isolated from the other IPSpaces.
- All IPSpace names on a storage system must be unique. However the IPSpace names on active/active configuration partners must be the same.

There can be a maximum of 101 IPSpaces per storage system. One IPSpace (Default IPSpace) is created by default when the MultiStore license is enabled, so you can create the remaining 100. However, considering the fact that one physical network interface can belong to only one IPSpace, dedicating at least one physical interface per IPSpace limits the number of IPSpaces that can be set up on a storage system to the number of physical interfaces that are available on the storage system. VLAN tagging can be used to overcome this limitation:

- Using VLAN tagging, more than one IPSpace can share the same physical network interface so that more IPSpaces can be set up.
- VLAN tagging with IPSpaces ensures that network packets are forwarded to the appropriate IPSpace and securely delivered to the vFiler unit that is contained inside the IPSpace.

Note: VLAN tagging can be used without IPSpaces. Additional IPSpaces other than the default are not required but are recommended for any situation that requires a unique routing table.

4.1.4 Network configuration example

The following steps illustrate, at a high level, the process of configuring the network of a vFiler unit in a non-default IPSpace. For a detailed description, see Appendix B, “Network configuration recommendations for Data Motion” on page 75.

To configure the network of a vFiler unit in a non-default IPSpace:

1. Configure VIF on the storage controller. In this example, we use a VIF named smvif36.
2. Create two VLANs (1 and 2) over the VIF:

   ```
   vlan create smvif36 1 2
   ```

3. Enable the VLANs in the corresponding network switches.
4. Create the IPSpace where the vFiler unit will reside using the Data ONTAP system CLI:

   ```
   ipspace create ipspace1
   ```

5. To assign an interface to the IPSpace, make sure that the interface does not have a configured IP address. In the following example, VLAN 1 interface over the VIF which will be assigned to the IPSpace is cleared in case it already has a configured IP address:

   ```
   ifconfig smvif36-1 0.0.0.0
   ```

6. Assign the interface to the IPSpace. In the following example, the single mode VIF is assigned to the IPSpace that we created in step 1:

   ```
   ipspace assign ipspace1 smvif36-1
   ```

7. Check the available IPSpaces and the corresponding interfaces:

   ```
   ipspace list
   ```

8. Create the vFiler unit in the IPSpace created:

   ```
   vfiler create vfiler_test -n -s ipspace1 -i 10.73.66.139
   /vol/vfiler_test
   ```

9. Configure the VIF interface that is used with the IP address of the vFiler unit:

   ```
   ifconfig smvif36-1 10.73.66.139 netmask 255.255.255.0
   ```

10. Add the default route to the vFiler unit:

    ```
    vfiler run vfiler_test route add default 10.73.66.1 1
    ```

11. To create another vFiler unit in the same IPSpace, follow the same procedure as step 5:

```
vfiler create vfiler_test_1 -n -s ipspace1 -i 10.73.66.173
/vol/vfiler_test_1
```

12. Add the IP address of the new vFiler unit as an alias to the VIF:

```
ifconfig smvif36-1 alias 10.73.66.173 netmask 255.255.255.0
```

13. Add the default route to the second vFiler unit:

```
vfiler run vfiler_test_1 route add default 10.73.66.1 1
```

Note: For the network configurations to persist across reboots, add the ifconfig and the route commands used in steps 9, 10, 12 and 13 to the /etc/rc file. Optionally, if the hosting storage system is part of a high-availability configuration, edit the /etc/rc file in each partner of the HA configuration to define a partner interface for each interface that the vFiler unit uses.

4.1.5 Network configuration scenarios for Data Motion

This section covers several configuration examples for Data Motion.

vFiler units using IP aliases

Consider a typical IP storage infrastructure supporting Data Motion (see Figure 4-2 on page 42). Each N series storage system is connected to each of two switches to provide switch-level redundancy in the environment. The storage array on the left contains two vFiler units.

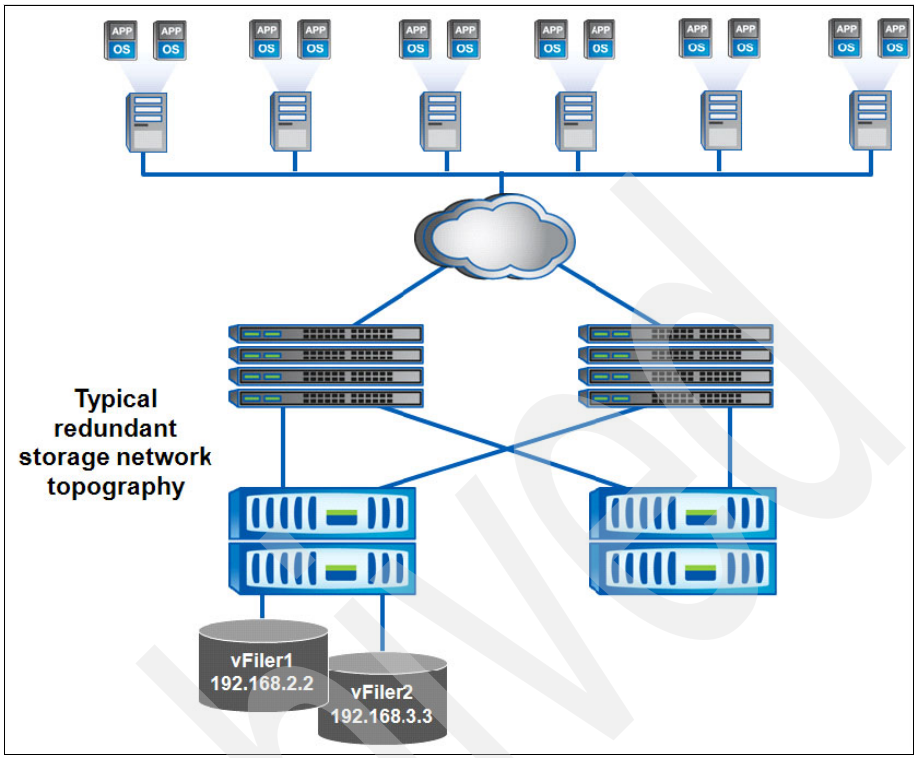

Figure 4-2 Typical redundant network topography

In Figure 4-3 on page 43, IP aliases are assigned to the vFiler units on the array. Each one is on a separate network—one on the 192.168.2.0 subnet, the other on the 192.168.3.0 subnet. Yellow lines represent the data paths that are used by the clients to access the vFiler units.

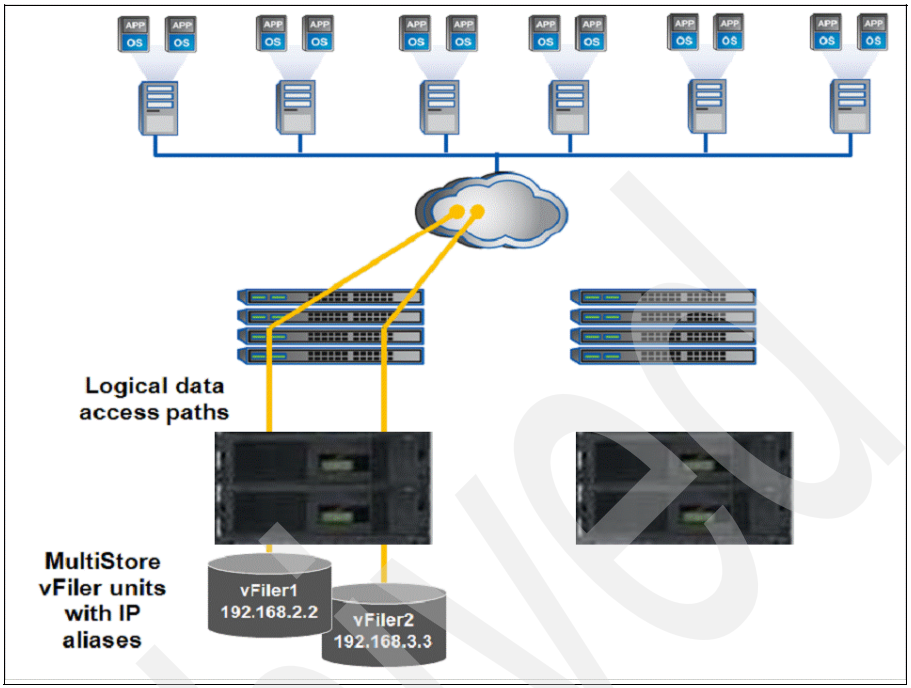

Figure 4-3 IP alias example

When a Data Motion migration is performed, the vFiler unit is migrated to the destination array and the IP alias is also created on the destination storage array, as shown in Figure 4-4 on page 44. The clients can access data in the new location. From the network's perspective, movement of the IP address is very similar to how an IP address moves when a link in a single-mode VIF fails and the IP address is moved to another port in the network.

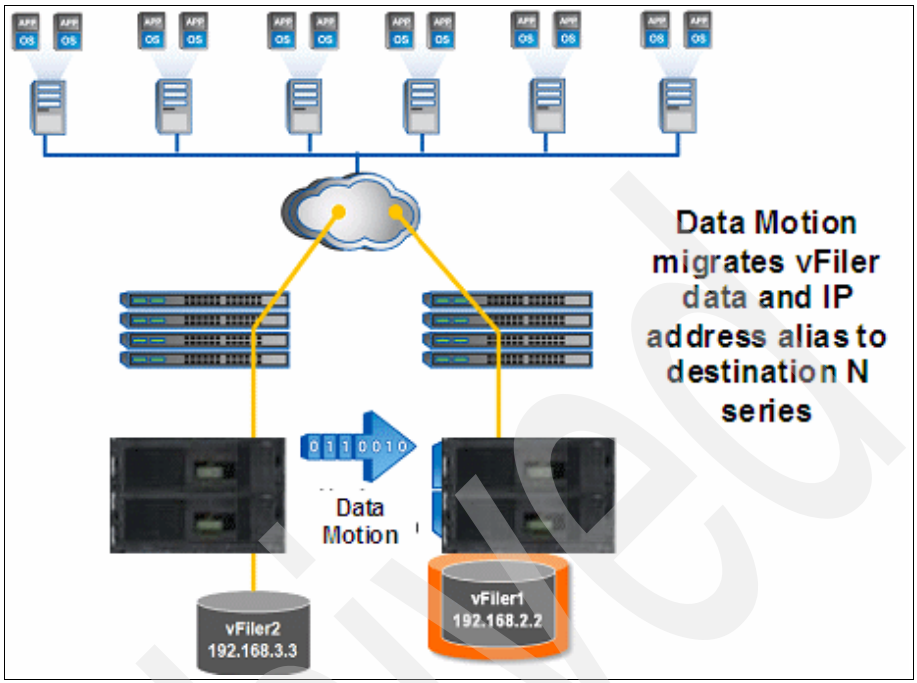

Figure 4-4 Data and IP migration

vFiler units using VLAN tagging and IPSpaces

In this example, the vFiler units are configured using IPSpaces and VLAN tagging on the storage controllers. vFiler1 is on subnet 192.168.2.0 VLAN. vFiler2 is on 192.168.3.0 VLAN. These VLANs might share the same physical ports in the network switch. This setup requires that the network switches have trunked the appropriate VLANs across all physical switch ports that might need access to those VLANs.

The setup in Figure 4-5 on page 45 is the recommended network configuration for Data Motion setup. Refer to Appendix B, “Network configuration recommendations for Data Motion” on page 75, for an end-to-end process of configuring the network for Data Motion. Appendix B, “” on page 75 covers network configuration best practices for network switches, storage systems, and vFiler units.

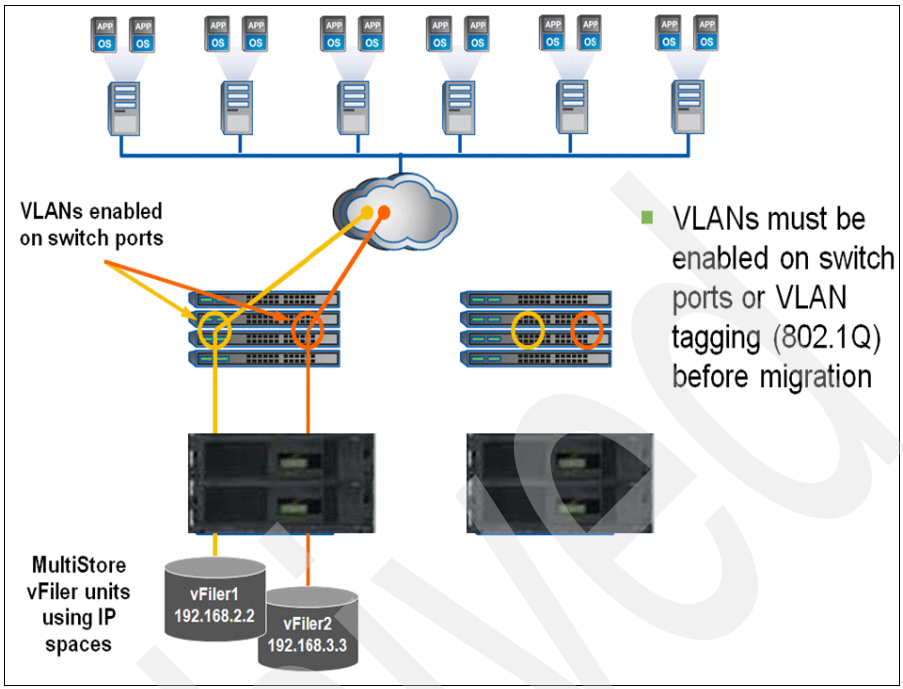

Figure 4-5 Recommended setup

Best practices

The best practices for vFiler units using VLAN tagging and IPSpaces are:

- Both 802.1q (VLAN tagging) and 802.3ad (Trunking) standards must be used together to enable VLANs to move without disruption.
- When using IPSpaces and VLAN tagging, take special attention to make sure that the team who manages the switches trunked the necessary VLANs across all switch ports on the redundant switches in the environment. Any physical switch port on the redundant switches through which the vFiler unit might be accessed must have the correct VLAN preconfigured or trunked, even if that port is never accessed by that VLAN prior to the Data Motion migration.

When migration is performed, Data Motion creates VLANs on the destination N series storage system, as shown in Figure 4-6 on page 46, but it cannot configure the VLAN trunking of ports in the network switch. Therefore, it is very important to configure this before migration. Data Motion cannot check the switch to make sure that the VLANs are properly configured. If the VLANs are not configured properly on the switch that is connected to the destination array, the

Data Motion migration will complete; however, the clients cannot access the vFiler unit at the new location, which can cause an application outage.

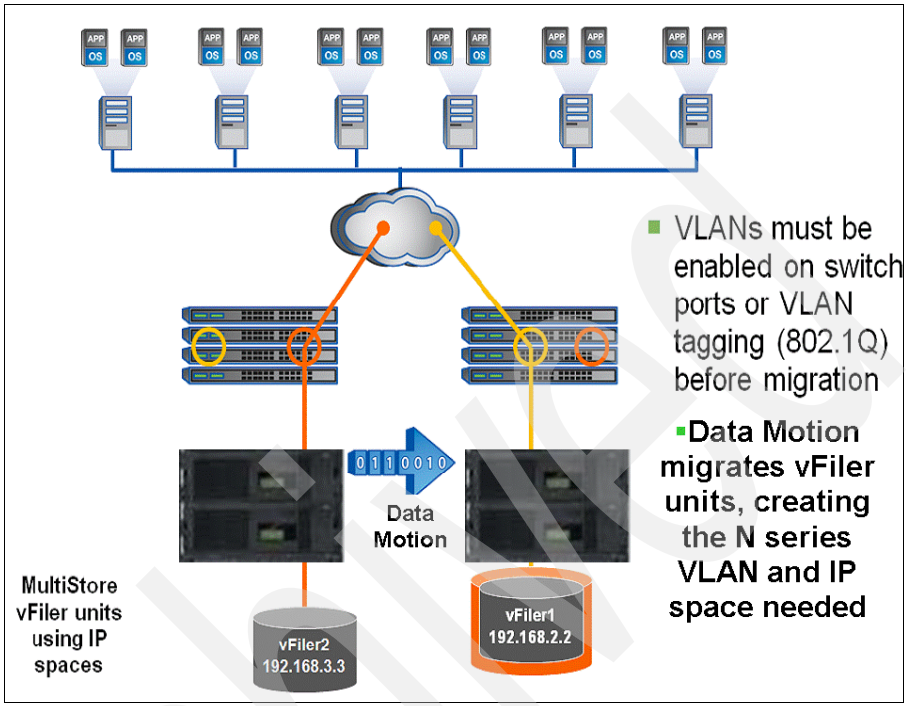

Figure 4-6 Data Motion creating VLANs on destination

4.2 Other networking considerations

Some of the other network considerations are:

- Before migration, make sure that the destination N series storage system is accessible by clients of the migrating vFiler unit:
  - VLANs must be pre-enabled on all possible destination switch ports.
- A bandwidth throttle feature is available for the baseline data transfer phase of Data Motion. This limit throttles SnapMirror streams for all of the FlexVols that are assigned to the vFiler unit, and Data Motion intelligently distributes this bandwidth across the vFiler volumes.

- Data Motion replicates the vFiler unit network configuration from source to destination, which includes jumbo frames support, IP spaces, routing tables, and vFiler options.
- Layered networking configurations are supported:
  - Multiple IP addresses per vFiler unit
  - Physical interfaces and multiple aliases
  - VLANs hosted on second-level VIFs
  - Static routes defined in the IP space of the vFiler unit
  - Shared IP spaces between multiple vFiler units

5

Data Motion best practices

In this chapter, we discuss the best practices that are associated with Data Motion. The topics that we discuss in this chapter are:

- Component-level best practices
- Process-level best practices
- Performance considerations
- IP network configuration best practices
- Storage-level best practices
- Best practices for supporting server virtualization environments

5.1 Component-level best practices

In this section, we discuss best practices for MultiStore, SnapMirror, and Provisioning Manager.

We recommend that you do not make any configuration changes until the Data Motion cutover phase completes. If any configuration change that is related to the vFiler unit occurs after starting the initial base line transfer, perform a Data Motion update before proceeding to the cutover phase.

Because the entire Data Motion process is controlled through Provisioning Manager, do not make any manual modifications to any configuration that is related to the vFiler units or SnapMirror (for example, the snapmirror.conf file).

Before starting the cutover process from Provisioning Manager, we recommend that you refresh the source and the destination storage systems in the Operations Manager using the **`dfm host discover`** command.

5.2 Process-level best practices

This section covers recommendations to consider as you execute Data Motion operations:

- If there is a long gap between initial baseline transfer phase completion and the subsequent cutover initiation, we recommend that you first perform an Update before proceeding to the cutover.
- Retain a stable backup of all storage entities (volumes and LUNs) that are associated with the vFiler before migration or cutover. In case something goes wrong, you can always revert to the stable state.
- The current release of Data Motion supports cutover or rollback only between same speed disk drives or from slower to faster speed disk drives. If you plan to do rapid rollback or rapid cutover types of operations after successful migration, we recommend that you carefully check the disk type of the source and destination aggregates between which the vFiler migrates. Rapid rollback or rapid retry after a successful migration is only possible between the aggregates that have the same speed disk drives.
- The current release of Data Motion supports cutover or rollback only between similar storage system models or from smaller to larger storage system models. If you plan to do rapid rollback or rapid cutover type of operations after successful migration, we recommend that you carefully check the model type of the source and destination storage systems between which the vFiler

migrates. Rapid rollback or rapid retry after a successful migration is only possible between the similar storage system models

5.3 Performance considerations

When the cutover phase is initiated, Provisioning Manager uses a built-in algorithm to decide whether the cutover can be successful. Various dynamic factors like CPU utilization in the destination and source storage systems, disk I/O on the aggregate of the destination system, and so on are taken into account for that decision-making process. If Provisioning Manager decides that the cutover cannot be completed successfully, it does not proceed with the cutover process. In such a case, the migration status is left at the cutover required state. In that scenario, Data Motion cutover needs to be reinitiated from Provisioning Manager.

We recommended that during Data Motion cutover, the CPU and disk I/O load on the source and destination storage systems must be in low-to-medium range. Because semi-sync SnapMirror is involved in the cutover phase, keep additional headroom for the CPU and disk I/O load before initiating cutover.

Tip: While choosing a Data Motion destination, make sure that it can handle the workload that is being migrated.

In this initial release of Data Motion, it is recommended that customers actively monitor the system performance during Data Motion cutover operation for systems that have dedupe or FlexClones at the Files or LUN level (on either the source or destination system). If excessive load is encountered during the migration, Data Motion can be aborted by the storage administrator.

Refer to “Support for deduplication and File/LUN FlexClone” on page 32 for deduplication and File/LUN FlexClone additional information and recommended monitoring methods during migration, where dedupe or FlexClones at the Files or LUN level exist on either the source or destination system.

5.4 IP network configuration best practices

In this section, we cover the IP best practices for Data Motion:

- We recommend that you use VLAN trunking on the network switches of any storage system that might be a source or destination for Data Motion.

- All required VLANs must be pre-enabled on destination switch ports. Before migration, make sure that the destination N series is accessible by clients of the migrating vFiler unit:
  - When the migration is performed, Data Motion creates the VLANs on the destination N series storage system, but it cannot configure the VLAN trunking of ports on the network switch. Therefore it is very important make these configurations before migration. Data Motion cannot check the switch to make sure that the VLANs are properly configured. If the VLANs are not configured properly on the switch that is connected to the destination array, the migration will complete; however, the clients will not be able to access the vFiler unit at the new location, and this can cause an application outage.
- The network of the vFiler unit must be configured inside an IP space that is associated with a VLAN interface over a VIF that is created on the storage system.
- We recommend that you use a separate IPSpace per vFiler unit (corresponding to various tenants, departments, or applications).
- For details, see 4.1, “IP requirements for Data Motion” on page 38.
- Appendix B describes a step-by-step networking best practices implementation.

5.5 Storage-level best practices

This section discusses several topics at the storage level and gives recommendations.

5.5.1 Required Data ONTAP configuration

The following Data ONTAP licenses must be installed for Data Motion: iSCSI and NFS, MultiStore, SnapMirror, Synchronous SnapMirror. For a complete list of hardware and software requirements, see 2.2, “Supported platforms and requirements” on page 15.

5.5.2 Support for deduplication and File/LUN FlexClone

Data Motion supports dedupe and FlexClones at the files or LUN level under low-to-medium load. To ensure the fastest migration times and hence minimal impact, it is recommended that Data Motion be performed in non-peak times or periods of lower load. See 5.5.2, “Support for deduplication and File/LUN

FlexClone” on page 52 for additional details and recommended monitoring during migration where dedupe or FlexClones at the files or LUN level exist on either the source or destination system.

5.5.3 Volume FlexClone support

FlexClone volumes can be members of vFiler units that are migrated with Data Motion. For details, see “Support for Volume FlexClone and LUN clone” on page 33.

5.5.4 Disk type support: FC, SATA, SAS

Data Motion supports migration between disks with the same speed or from slower to faster disks. Table 5-1 shows the supported configurations.

Table 5-1 Disk technology support

Source Aggregate disk type	Destination Aggregate disk type	Data Motion?
FC	FC	Allowed
SAS	SAS	Allowed
SATA	SATA	Allowed
ATA	ATA	Allowed
BSAS	BSAS	Allowed
FC	SAS	Allowed
SAS	FC	Allowed
ATA	SATA	Allowed
ATA	BSAS	Allowed
SATA	FC	Allowed

Source Aggregate disk type	Destination Aggregate disk type	Data Motion?
BSAS	FC	Allowed
SATA	SAS	Allowed
ATA	SAS	Allowed
BSAS	SAS	Allowed

5.5.5 SnapDrive and SnapManager best practices

When using SnapDrive, SnapManager, and Data Motion keep in mind the following recommendations:

- Do not create or delete any LUN inside volumes that are associated with the vFiler unit after starting the migration process and before cutover completes. We recommend that you perform all vFiler resource and dataset-related change operations before the initial baseline transfer.
- During the cutover or rollback, any backup and restore operations are unsuccessful for SnapManager products (SME, SMSQL, SMVI, SMHV, SMO, SMSAP, and SMOSS) because there are no retry options from SnapDrive (SDW and SDU). We recommend that you refrain from using any backup and restore commands during the cutover or rollback phases.

5.5.6 Host operating system disk time out settings

Observe the following timeout values when using Data Motion:

- Set the disk timeouts on the host operating systems using the N series Host Utilities Kit (HUK).
- Sample disk timeout values:
  - 190 sec for a guest operating system that is running on VMware ESX
  - 120 sec for a stand-alone Windows host without MPIO
  - 20 sec or 60 sec (for HUK v5.1 and after) with MPIO

5.5.7 Cluster failover

Data Motion does not support cutover during a clustered failover mode when a storage controller of an active-active HA pair takes over the other. If the takeover event happens after the cutover operation starts, the cutover will fail.

5.5.8 MetroCluster

Data Motion does not support cutover when one storage node of the MetroCluster system takes over the partner either automatically or manually (CFOD). If the takeover event occurs after the cutover operation starts, the cutover will fail.

5.6 Best practices for supporting server virtualization environments

This section covers recommendations for the virtualized environment and N series.

5.6.1 VMware

When using VMware ESX server with Data Motion, perform the following tasks:

1. Set the disk timeout values in the guest operating systems that are running on a VMware ESX server using the N series Host Utilities kit.
2. If Provisioning Manager is used to create and assign storage volumes (NFS) to vFiler units, those NFS volumes are exported as qtrees. When VMware ESX server data stores are created corresponding to these NFS exports, due to a limitation in Provisioning Manager, the storage savings that are obtained through various N series storage efficiency technologies on the data store are not directly visible in the VMware vCenter management interface. A workaround to overcome this limitation is to manually modify the exports that Provisioning Manager creates to export the root of the FlexVol volume (that is, the name of the data set) instead and mount that on the VMware ESX server. For VMware VI3 and vSphere, we recommend that you use FlexVol volumes instead of qtrees for VMware NFS data stores, which leads to a simplified architecture and also makes the value of N series storage efficiency technologies visible to the VMware server administrator.

A

The Data Motion process

In this appendix, we perform a migration of a vFiler unit that is serving live host and application data.

Tasks

The tasks that are involved with this migration are:

1. Create a vFiler unit and provision storage:
   - Method 1: Using N series CLI
   - Method 2: Using FilerView®
   - Method 3: Using Provisioning Manager
2. Assign the storage from the vFiler unit to the hosts and applications.
3. Perform online migration of the vFiler unit using Provisioning Manager.

Creating vFiler units and provisioning storage

There are three ways to create the vFiler unit assign storage. In this section, we discuss all three ways. Use one of these methods according to the environment and requirements.

Method 1: Using the N series storage system CLI

The N series storage system CLI offers a rich set of commands for creating and configuring vFiler units. You can perform the following management tasks from the CLI:

- Enable and disable the MultiStore license
- Allow and disallow protocols to be run on a vFiler unit
- Create a vFiler unit
- Set up a vFiler unit
- Start and stop a vFiler unit
- Destroy a vFiler unit
- Move resources to and from a vFiler unit
- Monitor the status of a vFiler unit

vFiler command set

The vFiler command sets are:

```
vfiler help - Help for vfiler command
vfiler context - Set the vfiler context of the CLI
vfiler create - Create a new vfiler
vfiler rename - Rename an existing vfiler
vfiler destroy - Release vfiler resources
vfiler dr - Configure a vfiler for disaster recovery
vfiler add - Add resources to a vfiler
```

```
vfiler remove - Remove resources from a vfilerSample caption
vfiler migrate - Migrate a vfiler from a remote filer
vfiler move - Move resources between vfilers
vfiler start - Restart a stopped vfiler
vfiler stop - Stop a running vfiler
vfiler status - Provide status on vfiler configuration
vfiler run - Run a command on a vfiler
vfiler allow - Allow use of a protocol on a vfiler
vfiler disallow - Disallow use of a protocol on a vfiler
vfiler limit - Limit the number of vfilers that can be created
```

Method 2: Using the FilerView interface

The FilerView graphical user interface (GUI) supports the creation and configuration of vFiler units. You can perform the following management tasks from the FilerView GUI:

- Enable and disable the MultiStore license
- Allow and disallow protocols to be run on a vFiler unit
- Create a vFiler unit
- Set up a vFiler unit
- Start and stop a vFiler unit
- Destroy a vFiler unit
- Move resources to and from a vFiler unit
- Monitor the status of a vFiler unit

Method 3: Using Provisioning Manager

Provisioning Manager automates the process of vFiler creation and storage provisioning.

Figure A-1 on page 60 depicts the high-level steps in the automated creation and provisioning of vFiler units by using Provisioning Manager.

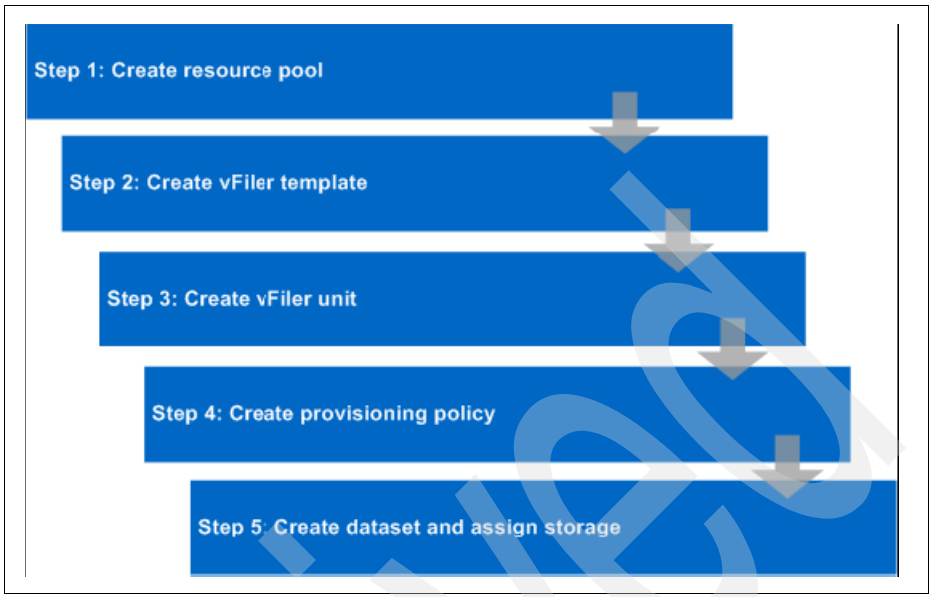

Figure A-1 high level provisioning steps

Assigning the storage from the vFiler unit to the hosts and applications

After the vFiler unit is created and configured, the storage (volumes and LUNs) that is assigned to the vFiler unit can be mapped to the hosts and the applications, similar to any N series storage system.

vFiler unit online migration using Provisioning Manager

Initiate the online migration of the vFiler unit using Data Motion from the Provisioning Manager.

Initial baseline transfer

To initiate an initial baseline transfer:

1. Select the vFiler unit and click Hosts → Start migration → vFiler Units, as shown in Figure A-2 on page 61.

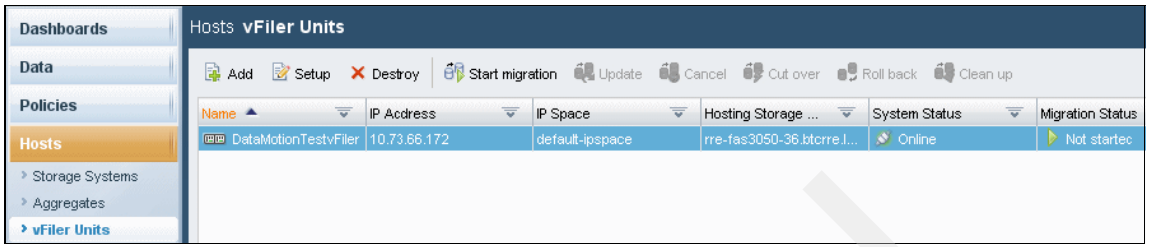

Figure A-2 Initial baseline transfer

The Welcome panel is displayed, as shown in Figure A-3.

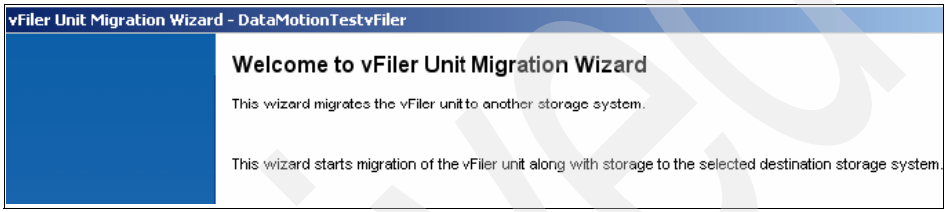

Figure A-3 Welcome panel

2. Select **Online migration**, as shown in Figure A-4.

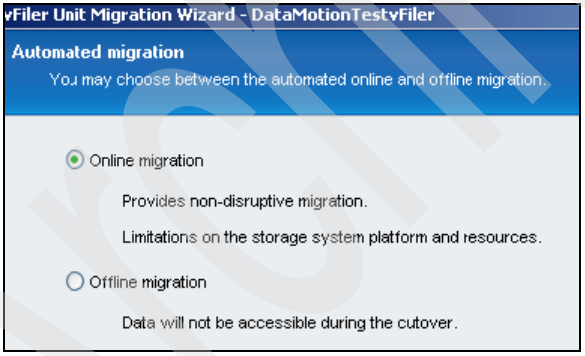

Figure A-4 Migration selection

3. Select the destination N series storage system to which the vFiler unit is to be migrated, as shown in Figure A-5 on page 62.

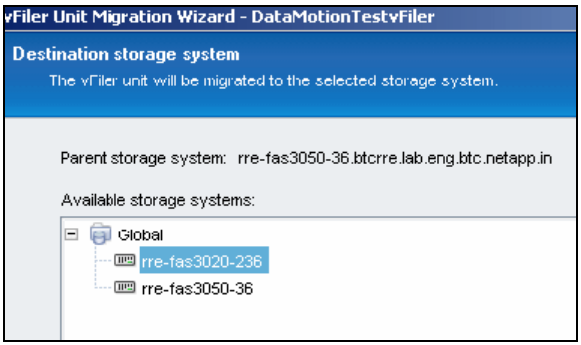

Figure A-5 Selecting N series

4. The network interface on the destination N series system must be configured carefully for online migration. Select the network interface to use and the VLAN ID, if any, as shown in Figure A-6.

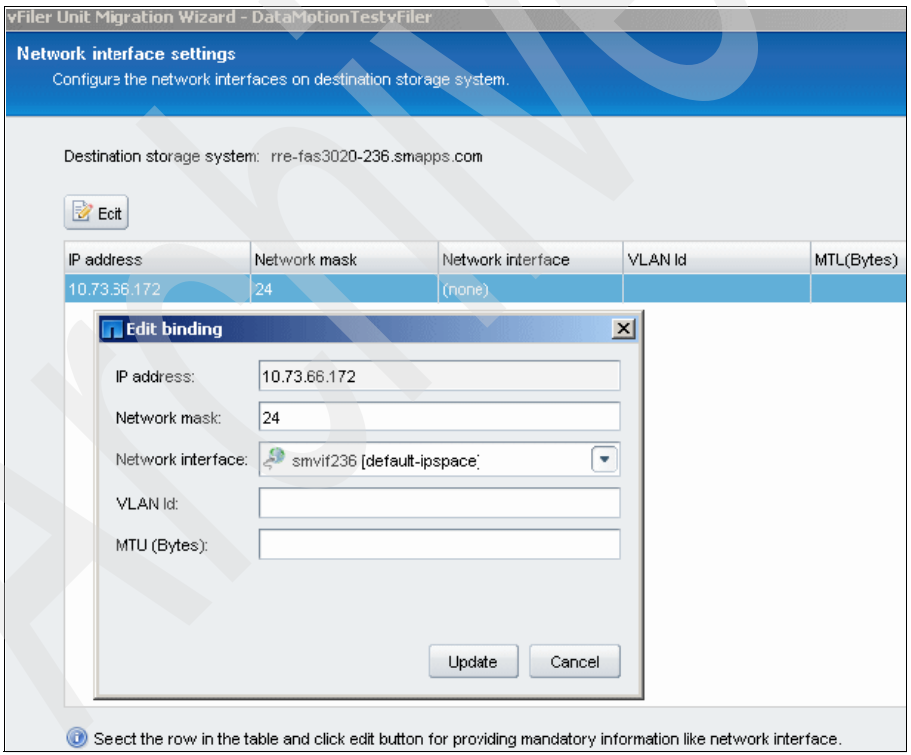

Figure A-6 Selecting the Network interface

5. (Optional) Throttle the bandwidth to be used for the initial baseline transfer, as shown in Figure A-7 on page 63.

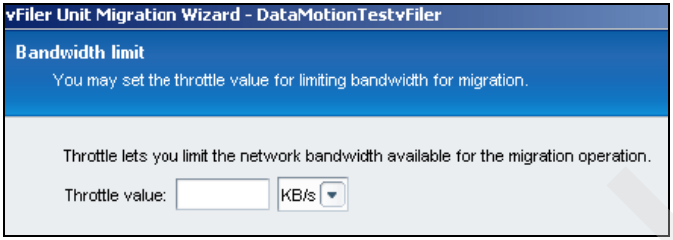

Figure A-7 Bandwidth throttle

6. There is also an optional feature to cutover to the destination N series automatically after initial data transfer. This feature provides the option of One Step Migration. For our example, Figure A-8, this option is not selected.

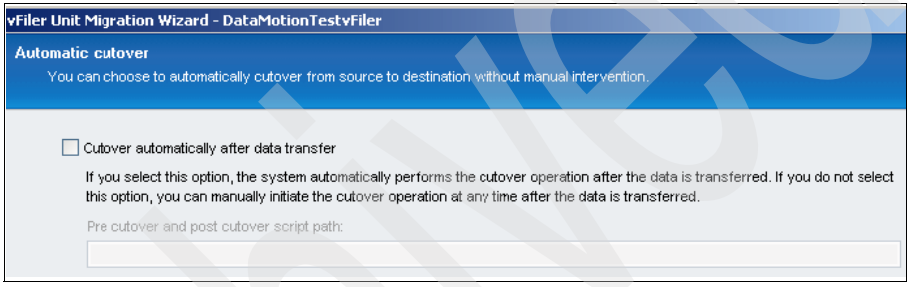

Figure A-8 Migration configuration

7. Preview the steps that are automated by Provisioning Manager for the vFiler migration process, as shown in Figure A-10 on page 64.

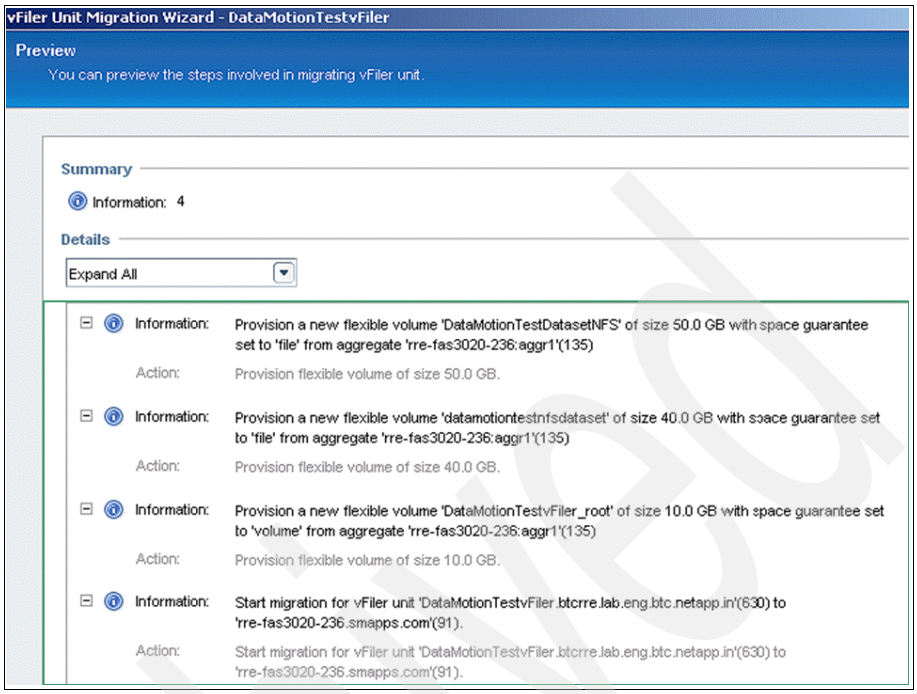

Figure A-9 Preview

8. Complete the vFiler Unit Migration Wizard. Figure A-10 shows the Completing Migration wizard.

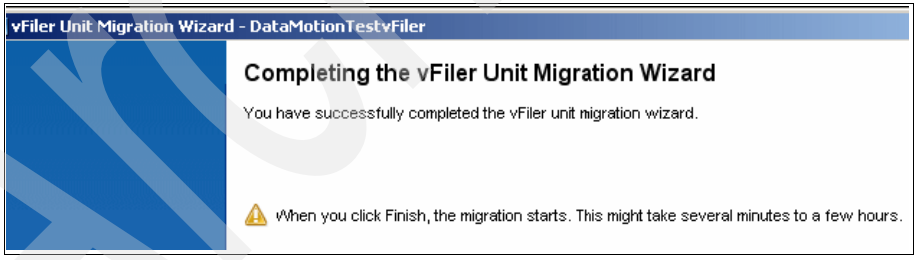

Figure A-10 Completing Migration wizard

A monitor window opens, displaying the status of the migration job, as shown in Figure A-11 on page 65.

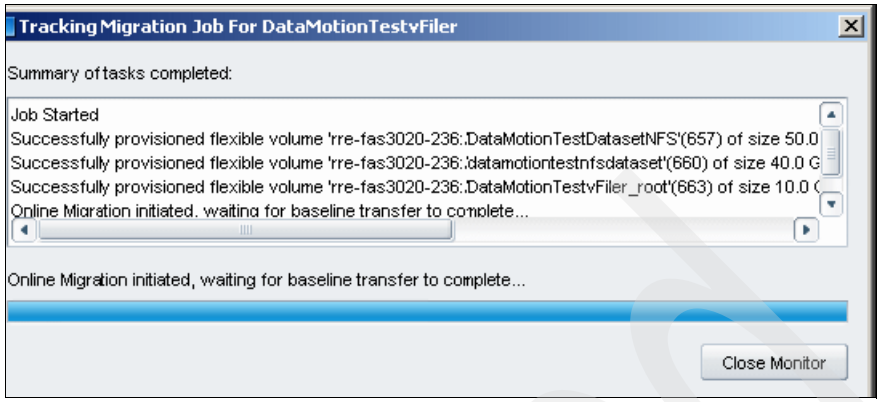

Figure A-11 Migration status

9. The status of the migration job can also be viewed from the Data Jobs panel, as shown in Figure A-12.

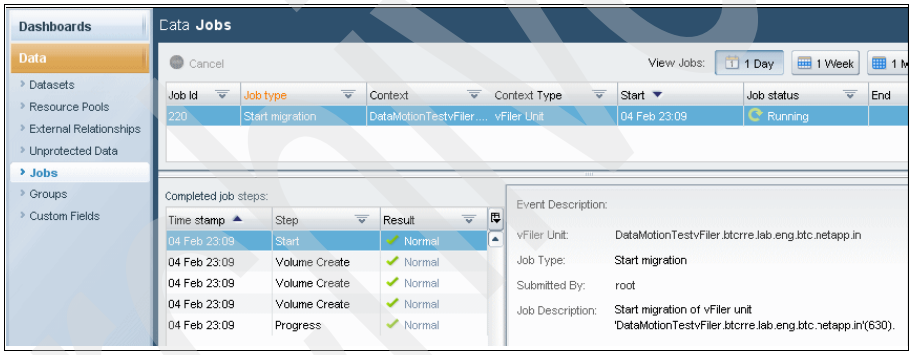

Figure A-12 Migration status

When the baseline transfer of the data is complete, the migration status of the vFiler unit changes to Started, Cutover Required. At this point, you have three options:

- Update the baseline created in the destination system (recommended if there is a long gap between the time when the baseline transfer completes and the time when the cutover process starts.
- Cancel the Data Motion process.
- Start the cutover process, shown in Figure A-13 on page 66, to migrate the vFiler unit to the destination system.

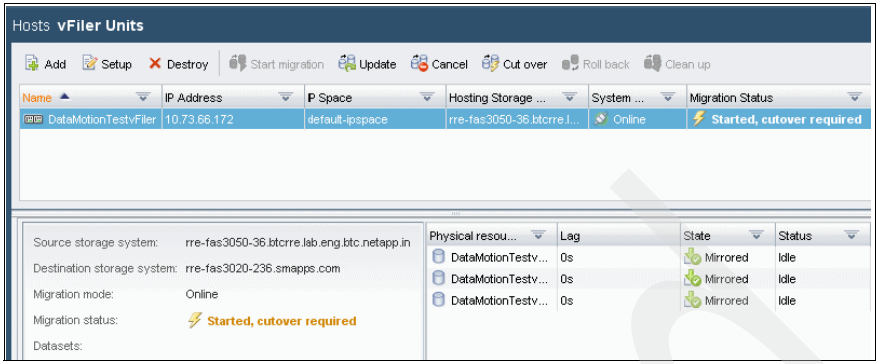

Figure A-13 starting cutover process

Cutover

To start the cutover process:

1. Select **Hosts** → **vFiler**, as shown in Figure A-14, and click **Cutover**.

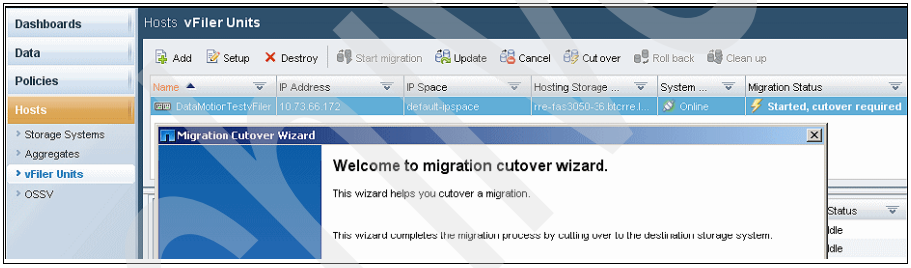

Figure A-14 Selecting vFiler

2. (Optional) Provide the path of the script to run before and after cutover.
3. Select an option for Network Routes Migration, as shown in Figure A-15 on page 67.

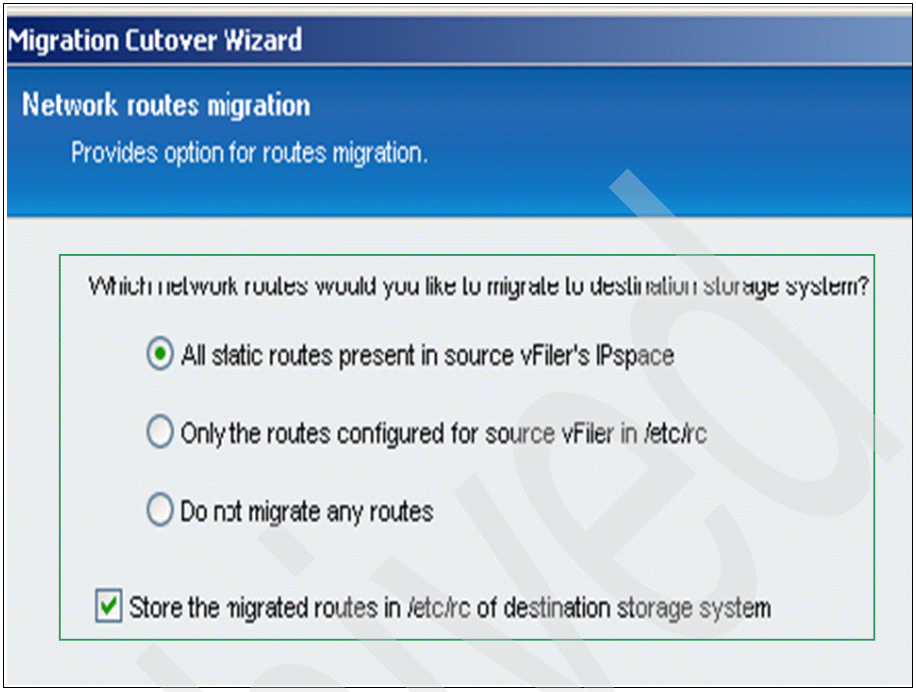

Figure A-15 Route selection

4. Review the summary of steps that Provisioning Manager automates for cutover to occur, as shown in Figure A-16 on page 68.

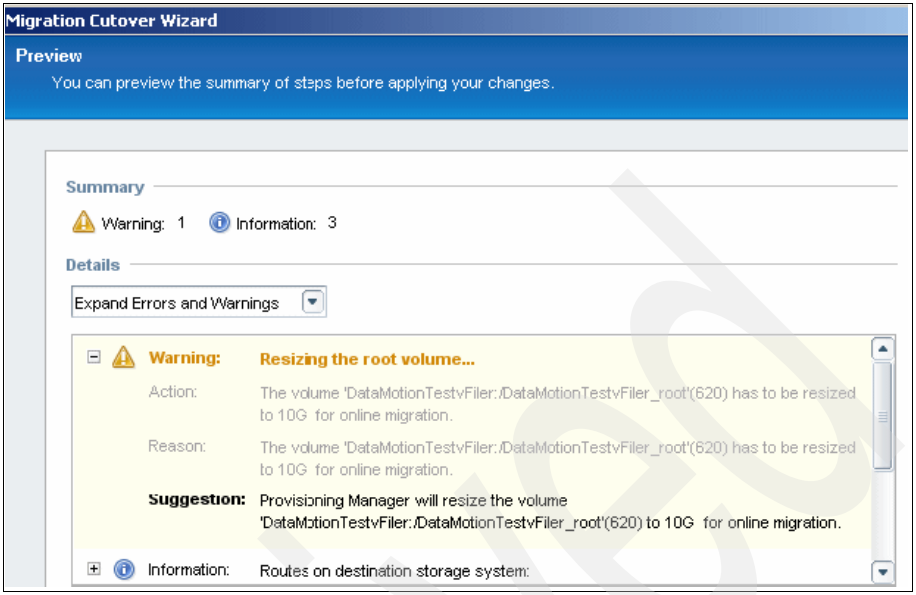

Figure A-16 Review

5. Complete the Migration Cutover Wizard.

   A monitoring window opens, displaying the status of the cutover job, as shown in Figure A-17.

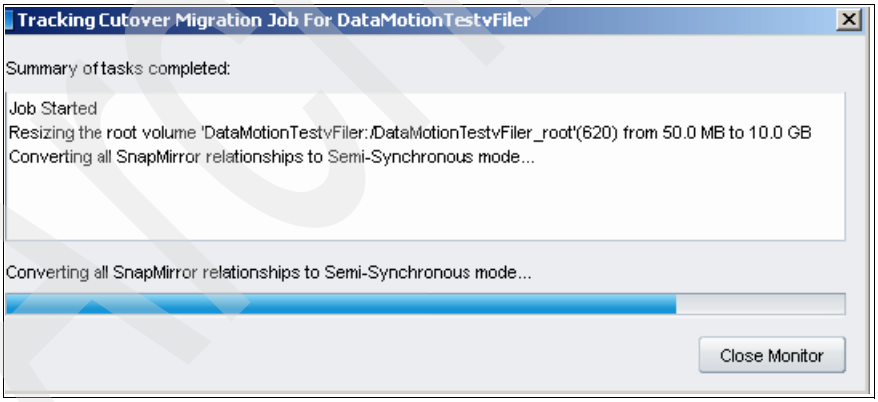

Figure A-17 Cutover status

6. The status of the cutover job can also be monitored from **Data** → **Jobs** → Job **Type: Cutover**, as shown in Figure A-18 on page 69.

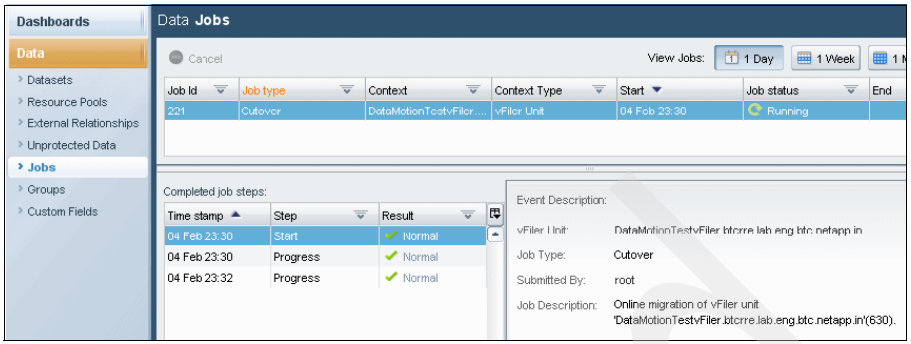

Figure A-18 Cutover status

When the cutover successfully completes, the migration status of the vFiler unit changes to Migrated, cleanup Required, as shown in Figure A-19. At this point, you can either Rollback the vFiler unit to the source storage system or Cleanup the vFiler unit resources from the source storage system.

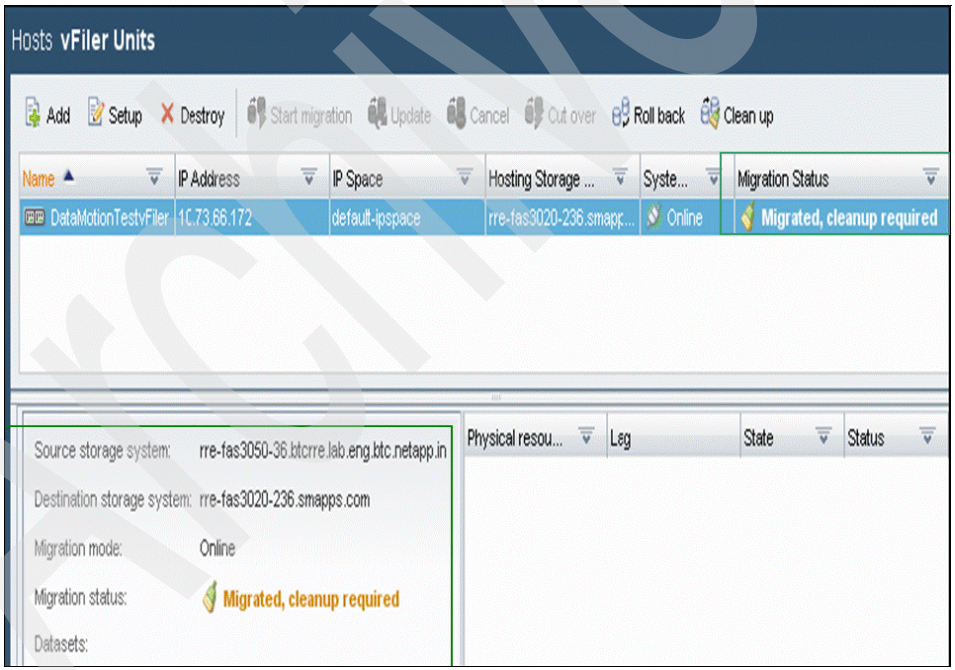

Figure A-19 vFiler options

Rollback

To rollback the vFiler unit to the source storage system:

1. Select **Hosts** → **vFiler Units** → **Rollback**, as shown in Figure A-20 on page 70.

Figure A-20 Rollback wizard

2. (Optional) Provide a path of the script that must run before and after rollback.
3. Select the required options for Network routes migration, as shown in Figure A-21.

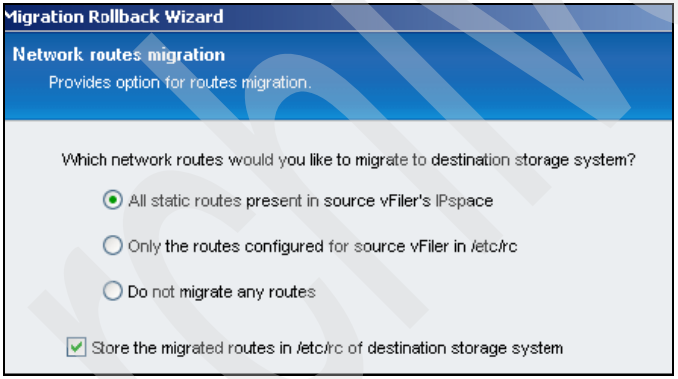

Figure A-21 Network routes

4. Review the steps that Provisioning Manager automates for Rollback, as shown in Figure A-22 on page 71.

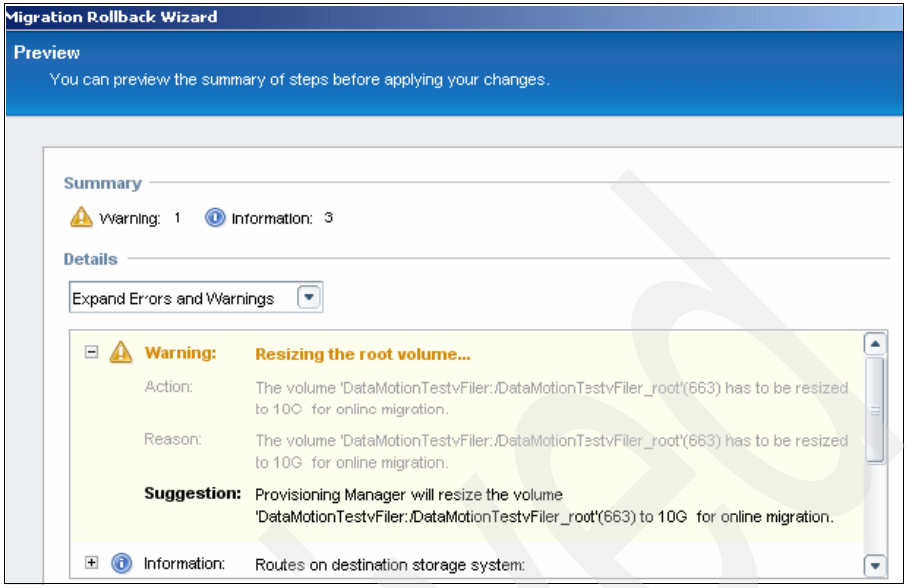

Figure A-22 Review the steps

5. Complete the Migration Rollback Wizard.

   A monitoring window opens, displaying the status of the rollback job, as shown in Figure A-23.

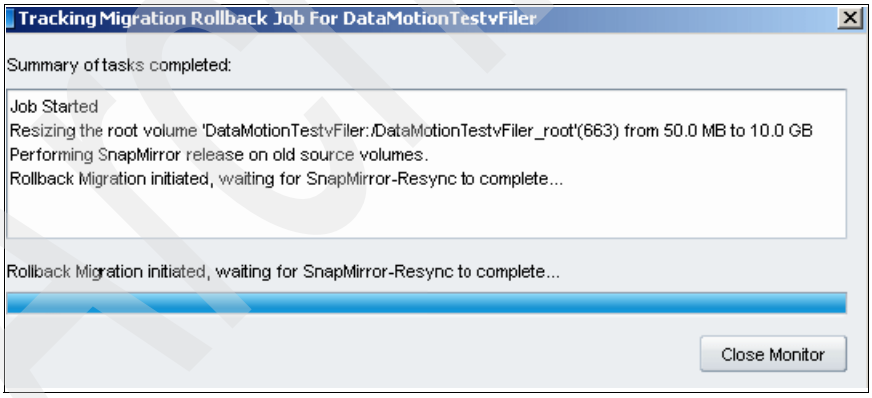

Figure A-23 Rollback status

6. The status of the Rollback job can also be monitored from **Data** → **Jobs** → **Job Type: Migration rollback**, as shown in Figure A-24 on page 72.

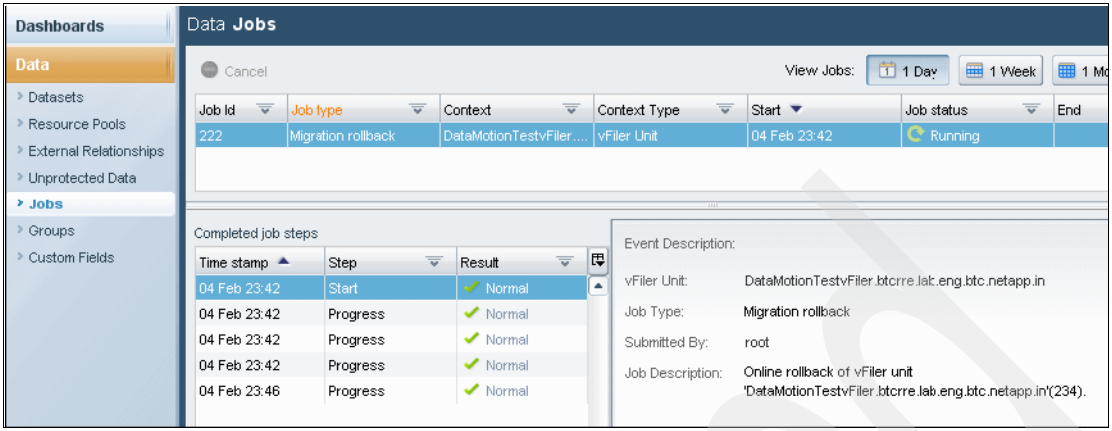

Figure A-24 Type of rollback

When rollback is completes successfully, the Migration Status of the vFiler unit changes to Rolled back. At this point, you can either cutover the vFiler unit to the destination storage system or cleanup the vFiler unit resources from the destination storage system, as shown in Figure A-25.

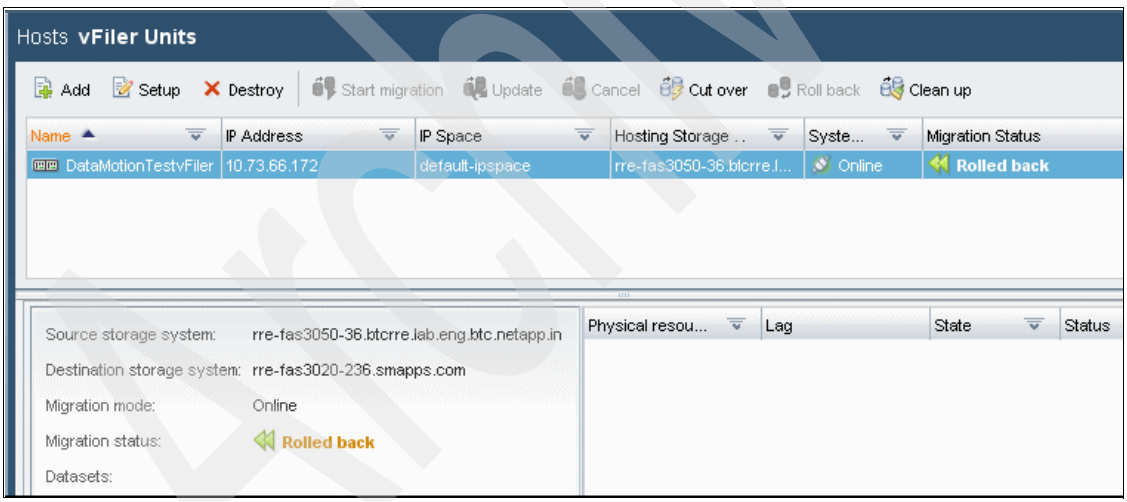

Figure A-25 vFiler options

Cleanup

To cleanup the vFiler unit and its resources from the source or destination storage system:

1. Select **Hosts** → **vFiler Units** → **Clean up**, as shown in Figure A-26 on page 73.

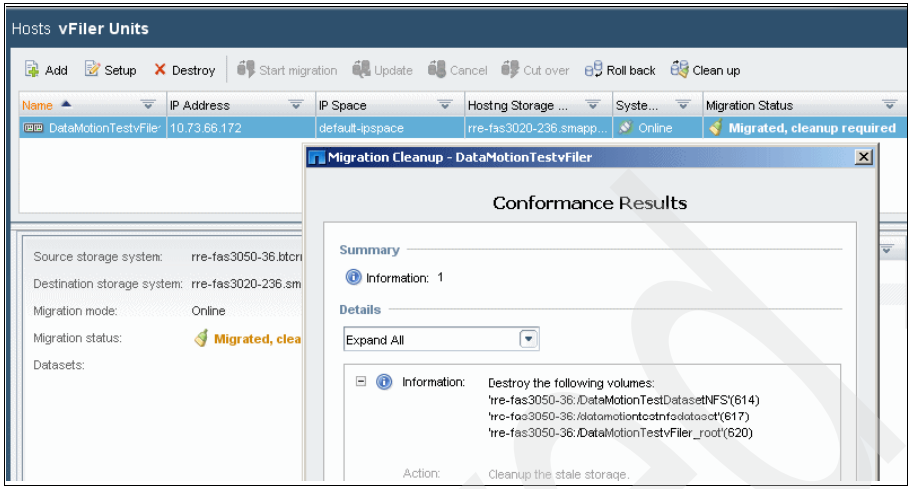

Figure A-26 Cleanup options

2. After cleanup, the storage system settings change, as shown in Figure A-27.

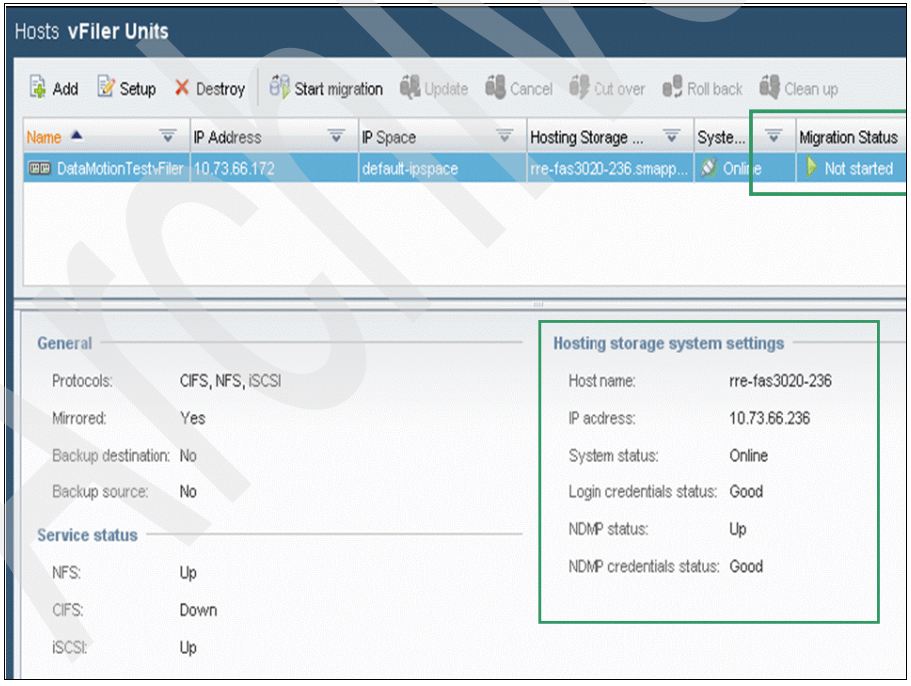

Figure A-27 Storage system settings display

B

Network configuration recommendations for Data Motion

Properly configuring the storage network is of utmost importance for successful and secure execution of Data Motion. In this section, we describe an example set up that is implemented in accordance with the Data Motion networking best practices.

Figure B-1 illustrates a sample set up. There are two high-availability pairs: FAS1 and FAS2 and FAS3 andFAS4. The vFiler unit called vFiler1 running on FAS1 must be migrated to FAS3 using Data Motion.

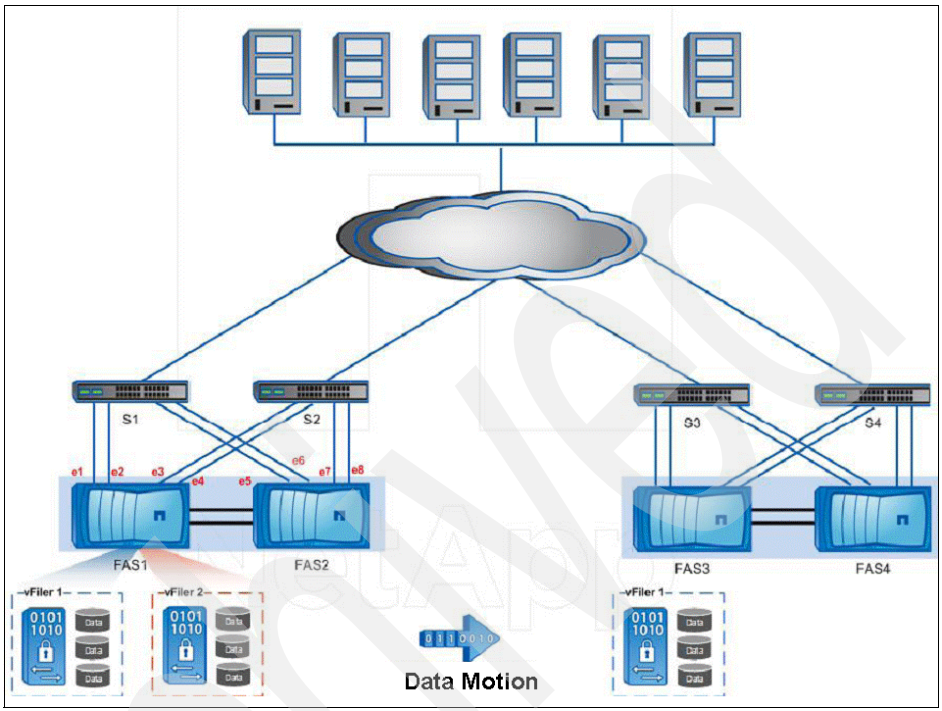

Figure B-1 Sample setup

The recommended method of configuring the networking on the storage systems and the network switches for Data Motion is covered in the following steps:

1. Configure VIF on the N series.
2. Create VLAN over the VIF.
3. Enable the VLAN in the corresponding network switches.
4. Create IPSpace.
5. Assign the appropriate VLAN interface over the VIF to the IPSpace.
6. Create the vFiler unit in the IPSpace with the associated VLAN interface over the VIF.
7. Assign an IP address to the appropriate VLAN interface over the VIF.
8. Add a default route to the vFiler.

Step 1: Configuring VIF on the storage controller

For simplicity, Figure B-2 shows the connection between just one N series storage controller and the network switches. However the same configuration must be repeated for all of the storage controllers and the corresponding switches.

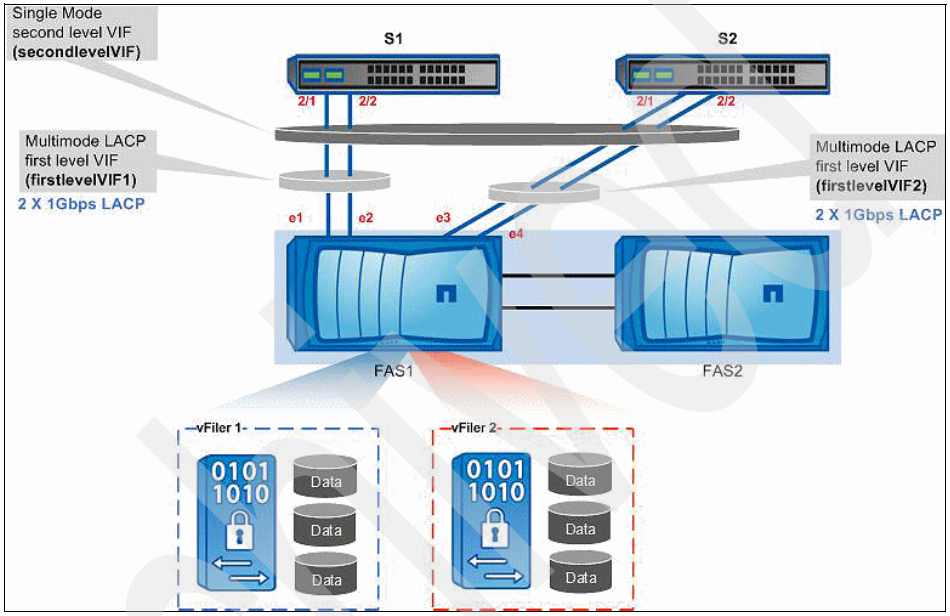

Figure B-2 N series connections

In this example, the four 1 Gbps links from the storage controller to two network switches are configured as a multilayered VIF. e1 and e2 form a first-level multimode LACP VIF, and e3 and e4 form another. These two first-level multimode LACP VIFs are then combined to form a second-level single-mode VIF:

```
vif create lacp firstlevelVIF1 –b ip e1 e2
vif create lacp firstlevelVIF2 –b ip e3 e4
vif create single secondlevelVIF firstlevelVIF1 firstlevelVIF2
vif favor firstlevelVIF1
```

Step 2: Creating VLAN over VIF

Two VLANs (1 and 2) are created over the second level VIF using the following command:

```
vlan create secondlevelVIF 1 2
```

Step 3: Enabling the VLANs in the corresponding network switches

Assuming that you are using Cisco network switches running IOS, configuration must be made in the switch to convert the switch ports from a standard access port to a VLAN trunked interface to support the VLAN defined.

In the switch S1, the configuration can look like this:

```
interface Port-channel1
description LACP Channel for NetApp e1-e2 2x 1Gbps LACP 802.1q Trunk
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1-2
switchport mode trunk
switchport nonegotiate
spanning-tree guard loop
spanning-tree portfast trunk
mtu 9216

interface GigabitEthernet 2/1
description NetApp e1 2x 1Gbps LACP 802.1q Trunk
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1-2
switchport mode trunk
switchport nonegotiate
spanning-tree guard loop
spanning-tree portfast trunk
channel-protocol lacp
channel-group 1 mode active
mtu 9216

interface GigabitEthernet 2/2
description Nseries e2 2x 1Gbps LACP 802.1q Trunk
switchport
```

```
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 1-2
switchport mode trunk
switchport nonegotiate
spanning-tree guard loop
spanning-tree portfast trunk
channel-protocol lacp
channel-group 1 mode active
mtu 9216
```

Step 4: Creating IPSpace

Create the IPSpace where the vFiler unit will reside:

```
ipspace create ipspace1
```

Step 5: Assigning the appropriate VLAN interface over the VIF to the IPSpace

To assign an interface to the IPSpace:

1. Make sure that the interface does not have a configured IP address. In the following example, the VLAN 1 interface over the second level VIF that will be assigned to the IPSpace is cleared in case it already has a configured IP address.

   ```
   ifconfig secondlevelVIF-1 0.0.0.0
   ```

2. Assign the interface to the IPSpace. In the following example, the single mode VIF is assigned to the IPSpace that we created in step 1:

   ```
   ipspace assign ipspace1 secondlevelVIF-1
   ```

3. Check the available IPSpaces and the corresponding interfaces:

   ```
   ipspace list
   ```

Step 6: Creating the vFiler unit

Create the vFiler unit in the IPSpace with the associated VLAN interface over the VIF. Create the vFiler unit in the IPSpace just created:

```
vfiler create vfiler1 -n -s ipspace1 -i 192.168.1.10 /vol/vfiler1
```

Step 7: Assigning an IP address to appropriate VLAN interface over VIF

Configure the VLAN interface over the VIF with the IP address of the vFiler unit:

```
ifconfig secondlevelVIF-1 192.168.1.10 netmask 255.255.255.0
```

Step 8: Adding the default route to the vFiler unit

To add the default route to the vFiler unit:

1. Add the default route to the vFiler unit:

   ```
   vfiler run vfiler1 route add default 192.168.1.1 1
   ```

   The setup looks like Figure B-3.

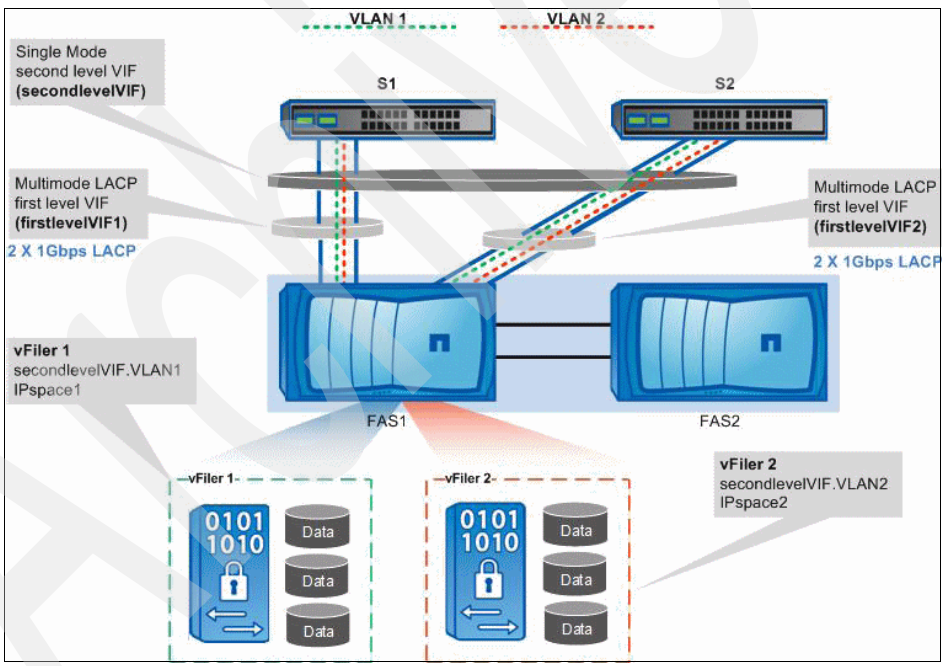

Figure B-3 Adding the default route to vFiler

Note: To create another vFiler unit in the same IPSpace, follow the same procedure as in step 6:

```
vfiler create vfiler_test -n -s ipspace1 -i 192.168.1.11
/vol/vfiler_test
```

Add the IP address of the new vFiler unit as an alias to the VLAN interface over VIF:

```
ifconfig secondlevelVIF-1 alias 192.168.1.11 netmask
255.255.255.0
```

2. Add the default route to the second vFiler unit:

```
vfiler run vfiler_test route add default 192.168.1.1 1
```

Note: To create another vFiler unit (vFiler2) in a different IPSpace, follow a procedure similar to steps 4 through 8 and associate the newly created IPSpace (ipspace2) with VLAN 2 over the VIF(secondlevelVIF-2):

```
ipspace create ipspace2
ifconfig secondlevelVIF-2 0.0.0.0
ipspace assign ipspace2 secondlevelVIF-2
vfiler create vfiler2 -n -s ipspace2 -i 192.168.2.10 /vol/vfiler2
ifconfig secondlevelVIF-2 192.168.2.10 netmask 255.255.255.0
vfiler run vfiler2 route add default 192.168.2.1
```

The sample setup described in this section uses a Cisco Catalyst switch with 2 x 1Gbps multimode VIFs to the first switch and the same to a second switch with single-mode VIF on top. Some other possibilities are:

- 2 X 10 Gbps single mode VIF for Cisco Catalyst switch
- 2 X 10 Gbps multimode VIF for Cisco Nexus switch running VPC

Steps to enable VLAN on the switch vary depending on the type of network switch that is selected. For a Cisco Nexus switch, here is a sample configuration:

```
interface Port-channel1
  description LACP Channel for NetApp e1-e2 2x 1Gbps LACP 802.1q
  Trunk
  switchport mode trunk
  switchport trunk allowed vlan 1,2
  spanning-tree port type edge trunk
  spanning-tree guard loop

interface Ethernet1/1
  description NetApp e1 2x 1Gbps LACP 802.1q Trunk
```

```
    switchport mode trunk
    switchport trunk allowed vlan 1,2
    channel-group 1 mode active
    spanning-tree port type edge trunk
    spanning-tree guard loop
    speed 10000
interface Ethernet1/2
    description NetApp e1 2x 1Gbps LACP 802.1q Trunk
    switchport mode trunk
    switchport trunk allowed vlan 1,2
    channel-group 1 mode active
    spanning-tree port type edge trunk
    spanning-tree guard loop
    speed 10000
```

Note: Cisco Nexus 5000 does not require an administrator to identify the VLAN trunk encapsulation type of 802.1q because it is enabled by default.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 84. Note that some of the documents referenced here might be available in softcopy only.

- *IBM System Storage N Series SnapMirror,* SG24-7260-00
- *Data Protection Strategies in IBM System Storage N Series*, SG24-7591-00
- *IBM System Storage N series with VMware ESX Server*, SG24-7636
- *IBM System Storage N series and VMware vSphere Storage Best Practices*, SG24-7871

Other publications

These publications are also relevant as further information sources:

- *Data ONTAP MultiStore Management Guide,* GC52-1281
- *IBM System Storage N series Provisioning Manager and Protection Manager Administration Guide for Use with DataFabric Manager, GC52-1288*
- *IBM System Storage N series Data ONTAP 7.3 Storage Management Guide,* GC52-1277-03

Online resources

These web sites and URLs are also relevant as further information sources:

- IBM System Storage N series and TotalStorage NAS interoperability matrixes

  http://www-03.ibm.com/systems/storage/nas/interophome.html

- Support for IBM System Storage and TotalStorage products

  http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/storageselectproduct?brandind=5000029&familyind=0&oldfamily=0Support for Data ONTAP
  http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?brandind=5000029&familyind=5329797&taskind=1

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

O

P

Q

R

S

T

V

IBM Redbooks

N series Data Motion

(0.2”spine)
0.17”<->0.473”
90<->249 pages

N series Data Motion

Data Motion in virtualized environments

Data Motion in physical environments

Dynamic Data Center enablement

As customers and service providers begin to consolidate more and more applications and workloads onto shared storage infrastructures, it becomes increasingly difficult to coordinate outages for planned downtime for things, such as hardware refreshes. This difficulty is because many users, groups, or customers might be using the shared storage infrastructure at the same time. Users expect these infrastructures to be available 24/7, so it is imperative that service outages that are required for storage life-cycle management, cost/service-level optimization, and any other planned downtime do not disturb the availability of the always-on infrastructure.

In this IBM Redbooks publication, we introduce you to the business value of Data Motion and its features. Storage and system administrators, data center managers, and IT as a service (ITaaS) providers will benefit from reading this book.

Data Motion significantly improves the availability of shared storage infrastructure by avoiding the service outages that are associated with planned activities, such as storage life-cycle management and cost/service-level optimization, thus helping you to enable an always-on IT environment.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7900-00 ISBN 0738434639