



Technical Report

## NetApp Data Motion

Larry Touchette, Roger Weeks, and Preetom Goswami, NetApp  
March 2010 | TR-3814

### **EXECUTIVE SUMMARY**

This document provides an overview of NetApp® Data Motion™, which is one of the foundation pieces of 'Always-On Data Mobility', providing non-disruptive data migration. It also describes the best practices for using NetApp Data Motion.

## TABLE OF CONTENTS

<b>1</b>	<b>INTRODUCTION</b>	<b>4</b>
1.1	PURPOSE	4
1.2	SCOPE	4
1.3	INTENDED AUDIENCE	4
1.4	PREREQUISITES	4
1.5	REFERENCES	4
<b>2</b>	<b>OVERVIEW OF NETAPP DATA MOTION</b>	<b>5</b>
2.1	BUSINESS CHALLENGE	5
2.2	WHAT IS DATA MOTION	5
2.3	BUSINESS VALUE OF NETAPP DATA MOTION	5
<b>3</b>	<b>NETAPP DATA MOTION USE CASES</b>	<b>6</b>
3.1	DATA MOTION IN VIRTUALIZED ENVIRONMENTS	6
3.2	NETAPP DATA MOTION IN PHYSICAL ENVIRONMENTS	7
3.3	ENABLING THE DYNAMIC DATA CENTER	7
<b>4</b>	<b>COMPONENTS OF NETAPP DATA MOTION</b>	<b>8</b>
4.1	MULTISTORE	8
4.2	NETAPP PROVISIONING MANAGER	9
4.3	NETAPP SNAPMIRROR—ASYNCHRONOUS AND SYNCHRONOUS	9
<b>5</b>	<b>SUPPORTED PLATFORMS AND REQUIREMENTS</b>	<b>10</b>
5.1	SOFTWARE REQUIREMENTS	10
5.2	HARDWARE REQUIREMENTS	10
<b>6</b>	<b>NETAPP DATA MOTION WORKFLOW AND PROCESS</b>	<b>10</b>
6.1	PHASES OF NETAPP DATA MOTION	11
<b>7</b>	<b>SUPPORTED CONFIGURATIONS</b>	<b>16</b>
7.1	SUPPORTED MIGRATION SCENARIOS	16
7.2	SUPPORTED VFILER CONFIGURATIONS	18
7.3	FLEXVOL REQUIREMENTS FOR NETAPP DATA MOTION	19
7.4	USING NETAPP DATA MOTION WITH OTHER NETAPP PRODUCTS	20
<b>8</b>	<b>IP NETWORKING REQUIREMENTS FOR NETAPP DATA MOTION</b>	<b>22</b>
8.1	NETWORK CONFIGURATION EXAMPLE	23
8.2	NETWORK CONFIGURATION SCENARIOS FOR NETAPP DATA MOTION	24
8.3	OTHER NETWORKING CONSIDERATIONS	27
<b>9</b>	<b>NETAPP DATA MOTION BEST PRACTICES</b>	<b>28</b>
9.1	COMPONENT-LEVEL BEST PRACTICES: MULTISTORE, SNAPMIRROR, PROVISIONING MANAGER	28

9.2	PROCESS-LEVEL BEST PRACTICES.....	28
9.3	PERFORMANCE CONSIDERATION .....	28
9.4	IP NETWORK CONFIGURATION BEST PRACTICES .....	29
9.5	STORAGE-LEVEL BEST PRACTICES.....	29
9.6	BEST PRACTICES FOR SUPPORTING SERVER VIRTUALIZATION ENVIRONMENT .....	31
<b>APPENDIX A: THE NETAPP DATA MOTION PROCESS .....</b>		<b>32</b>
	OBJECTIVE OF THIS EXERCISE .....	32
	TASKS.....	32
	CREATE VFILER UNIT AND PROVISION STORAGE .....	32
	ASSIGN THE STORAGE FROM THE VFILER UNIT TO THE HOSTS AND APPLICATIONS .....	34
	VFILER UNIT ONLINE MIGRATION USING PROVISIONING MANAGER.....	34
<b>APPENDIX B: NETWORK CONFIGURATION BEST PRACTICES FOR DATA MOTION .....</b>		<b>46</b>
<b>REFERENCES .....</b>		<b>51</b>
<b>AUTHORS 51</b>		
<b>ACKNOWLEDGEMENT .....</b>		<b>51</b>

# 1 INTRODUCTION

## 1.1 PURPOSE

This document provides an overview and describes best practices for NetApp Data Motion.

## 1.2 SCOPE

This document covers the following topics:

- What is NetApp 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

This document does not replace any other papers or guides related to various components of NetApp Data Motion—MultiStore®, SnapMirror®, and Provisioning Manager.

## 1.3 INTENDED AUDIENCE

- Storage and system administrators
- Data center managers
- IT as a service (ITaaS) providers

## 1.4 PREREQUISITES

This document assumes familiarity with the following:

- Basic knowledge of NetApp FAS systems and Data ONTAP®
- Overview of NetApp SnapMirror, MultiStore, and Provisioning Manager.

A complete understanding of all Provisioning Manager features is not necessary. Basic knowledge of host and vFiler™ management from Provisioning Manager is sufficient.

## 1.5 REFERENCES

Table 1) References.

Title	Location
Storage Virtualization and DR Using MultiStore® (vFiler™)	<a href="http://media.netapp.com/documents/tr-3462.pdf">http://media.netapp.com/documents/tr-3462.pdf</a>
MultiStore Management Guide	<a href="http://now.netapp.com/NOW/knowledge/docs/ontap/rel733rc1/pdfs/ontap/vfiler.pdf">http://now.netapp.com/NOW/knowledge/docs/ontap/rel733rc1/pdfs/ontap/vfiler.pdf</a>
Provisioning Manager 4.0 Workflow Guide	Product documentation in <a href="http://now.netapp.com">http://now.netapp.com</a>
Non-disruptive Operations Community	<a href="http://communities.netapp.com/groups/non-disruptive-operations">http://communities.netapp.com/groups/non-disruptive-operations</a>

## 2 OVERVIEW OF NETAPP DATA MOTION

### 2.1 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 activities like hardware refreshes. This is because many different users, groups, or customers may 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 required for storage lifecycle management, cost/service-level optimization, and any other planned downtime don't disturb the availability of the "always-on" infrastructure.

### 2.2 WHAT IS DATA MOTION

NetApp Data Motion is one of the foundation pieces of "always-on" data mobility, providing nondisruptive migration of MultiStore vFiler units for NetApp FAS, MetroCluster, and V-Series storage systems. It is the data mobility component of NetApp's cloud solutions.

Data Motion is capable of moving data from one Data ONTAP system to another while applications continue to run uninterrupted. 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, and/or application.

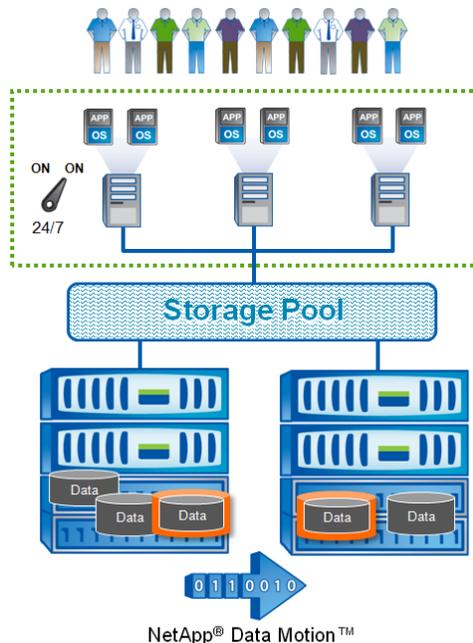


Figure 1) Data Motion.

### 2.3 BUSINESS VALUE OF NETAPP DATA MOTION

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

Here are some of the business values of Data Motion:

- 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

### 3 NETAPP DATA MOTION USE CASES

#### 3.1 DATA MOTION IN VIRTUALIZED ENVIRONMENTS

Virtualization is driving large-scale consolidation of many applications in the same shared infrastructure, resulting in significant costs saving through greater utilization, reduced infrastructure, minimized floor space, and smaller power needs. It is now much more difficult to coordinate outages across many different 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 like VMware® vSphere™ can provide nondisruptive data mobility at every layer of the virtual infrastructure.

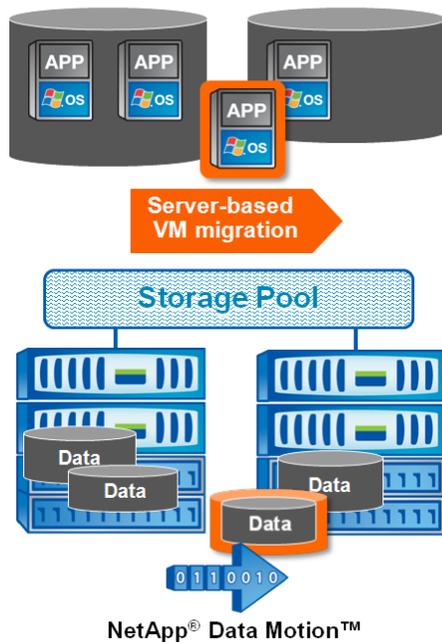


Figure 2) Data Motion in virtualized environment.

Virtual server software like VMware vSphere supports live migration of virtual machines across different server nodes in the cluster (VMotion®), and also provides the facility to move the virtual disks associated with the virtual machines online from one storage repository (or data store) to another (Storage VMotion). Data Motion complements migration technologies available from server virtualization vendors. It is important to recognize that the online migration technologies provided by Data Motion are designed to address the need for data mobility at a different 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 nondisruptively, for the reasons discussed in [section 2.3](#), 'Business Value of NetApp Data Motion'. Enterprise storage administration teams typically perform these operations.

Table 2 describes the online migration techniques that can be used in a virtual infrastructure.

Table 2) 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 associated with a VM	Virtual admin	Daily or weekly	Any customer using a virtual infrastructure
Storage infrastructure management (migration of entire datastore)	<b>NetApp Data Motion</b>	Movement of an entire datastore, customer, department, or application	Storage admin	Monthly	Enterprise, service provider, ITaaS, PaaS

### 3.2 NETAPP DATA MOTION IN PHYSICAL ENVIRONMENTS

#### DATABASE AND MESSAGING

Databases and messaging systems are the heartbeat of most companies, with nearly every staff member using them on 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 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 may go elsewhere. Data Motion removes the need for outages for lifecycle management, maintenance, and capacity and performance management of Web application environments.

### 3.3 ENABLING THE DYNAMIC DATA CENTER

NetApp Data Motion can be considered as one of the enablers of NetApp's vision of the dynamic data center, where avoiding planned downtime in any part of the infrastructure is of paramount interest. The following subsections describe some 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 provides mobility at the level of individual customers, departments, and/or application data sets that need to be isolated and secure.

#### TECHNOLOGY REFRESH

For simplified lifecycle management, new storage technology changes can be accomplished without the need for planned downtime. Installation of new storage controllers and higher capacity disk drives or shelves can all be done nondisruptively.

#### DYNAMIC CAPACITY SCALING

Where application data exceeds the storage capacity of an existing FAS or V-Series 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 SLA management, Data Motion can be used to move data across storage controllers for load balancing. Data can also be moved across storage tiers (slower disks to faster disks) to optimize performance.

## 4 COMPONENTS OF NETAPP DATA MOTION

NetApp Data Motion is a solution enabled by the combination of features from three NetApp products:

- [MultiStore](#)  
Provides the capability to partition a physical NetApp storage system into virtual systems called vFiler units, which is the unit of migration in NetApp Data Motion. In other words, MultiStore enables secure logical partitioning of network and storage resources in NetApp Data ONTAP into vFiler units.
- [Provisioning Manager](#)  
Provides the software interface for performing Data Motion migrations.
- [SnapMirror](#)  
Data is migrated by using NetApp SnapMirror. SnapMirror asynchronous mode is used for the initial baseline transfer phase of Data Motion. However, once 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 semi-synchronous mode of transfer.

### 4.1 MULTISTORE

NetApp MultiStore software provides separate and completely private logical partitions on a single NetApp 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.

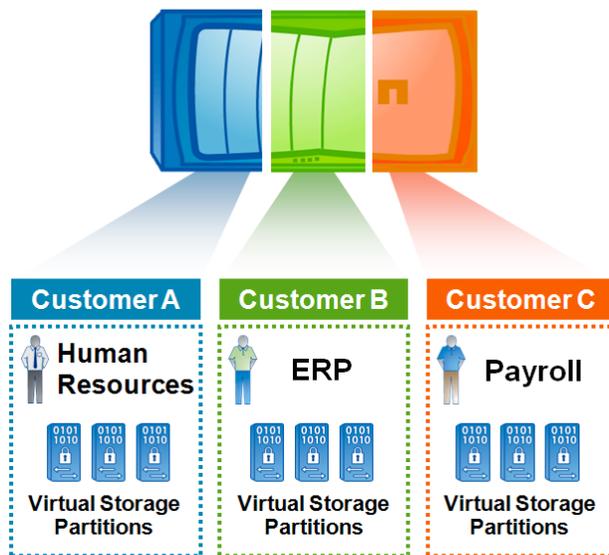


Figure 3) NetApp MultiStore.

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

For more information on MultiStore, see <http://media.netapp.com/documents/tr-3462.pdf>

## 4.2 NETAPP PROVISIONING MANAGER

Provisioning Manager is a feature in NetApp 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 due to incorrect storage configuration.

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

For more information about NetApp Provisioning Manager, see the product documentation in <http://now.netapp.com>

## 4.3 NETAPP SNAPMIRROR—ASYNCHRONOUS AND SYNCHRONOUS

SnapMirror is a data replication feature that efficiently mirrors data between two NetApp storage controllers. 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 RPO (recovery point objective—the amount of data not yet backed up) is dramatically lower than with legacy, 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 NetApp destination system is more than a simple backup or replica depot—it can also serve as the source system in a disaster recovery scenario. This means that the RTO (recovery time objective—the time from disaster to being back online) is also minimal.

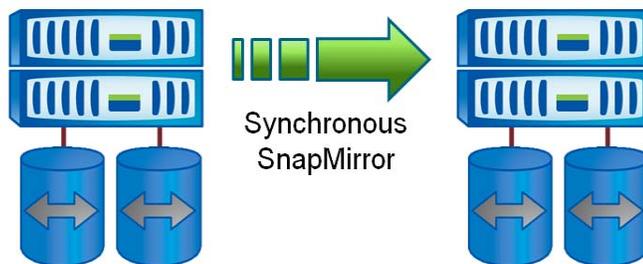


Figure 4) NetApp SnapMirror.

NetApp Data Motion transfers are initially set up in SnapMirror asynchronous mode; however, once 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<sup>®</sup> volume owned by a vFiler unit must be migrated in semi-sync mode at the same time.

For more information about NetApp SnapMirror, see:

<http://www.netapp.com/us/library/technical-reports/tr-3446.html>

<http://www.netapp.com/us/library/technical-reports/tr-3326.html>

## 5 SUPPORTED PLATFORMS AND REQUIREMENTS

### 5.1 SOFTWARE REQUIREMENTS

- Data ONTAP 7.3.3
  - Contains vFiler migration enhancements and APIs required by NetApp Data Motion
- 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

### 5.2 HARDWARE REQUIREMENTS

- NetApp FAS 20X0A, FAS 30X0A, FAS 31X0A, FAS 60X0A
- NetApp MetroCluster platforms
- NetApp V-Series systems—V6000, V3100

## 6 NETAPP DATA MOTION WORKFLOW AND PROCESS

NetApp Provisioning Manager provides the user interface for the NetApp Data Motion process and also controls the entire workflow. For an end-to-end Data Motion scenario, see [Appendix A](#).

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 preexisting vFiler units in the environment or vFiler units that were created with other tools or processes.

### ADVANTAGES OF USING PROVISIONING MANAGER FOR NETAPP DATA MOTION WORKFLOW CONTROL

- **Enforcement:** Provisioning Manager automatically enforces NetApp 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, postmigration cleanup, and so on—is not necessary. Some example scenarios that may disrupt a manual online migration process, which can be avoided by using Provisioning Manager are as follows:
  - Too much load on the systems involved
  - Configuration changes made between the SnapMirror initialize and the cutover (Example: new networks or volumes added)
  - Other SnapMirror relationships performing transfers and using transfer queues

## 6.1 PHASES OF NETAPP 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 is to provide 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.
- 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 NetApp 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, users need to manually initiate the cutover operation at any time after the data is transferred.

### INITIAL BASELINE TRANSFER PHASE

In this baseline transfer phase:

1. Initiate the Data Motion process from the Provisioning Manager interface.
2. Provisioning Manager performs the following steps:
  - i. Verifies that systems meet Data Motion requirements.
  - ii. Creates destination VLAN interfaces and IP spaces.  
(VLANs must be enabled on destination switch ports by the administrator.)
  - iii. Creates destination FlexVol volumes.
  - iv. Initiates SnapMirror baseline transfers.
  - v. SnapMirror schedule is set to asynchronously update every 3 minutes.
  - vi. When baseline completes, migration status is set to Cutover Required.

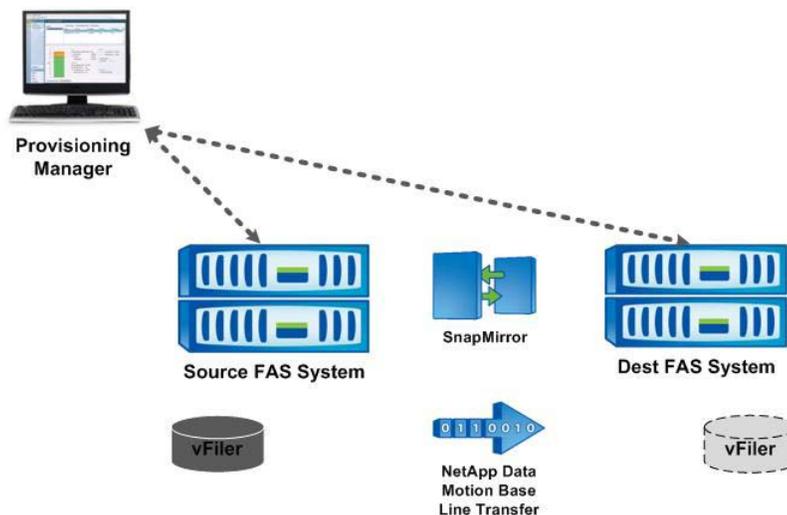


Figure 5) Baseline transfer.

Once the initial baseline transfer is complete and the migration status is changed to '**Started, cutover required**', user has the option perform an '**Update**' before proceeding to the next state (**Cutover**). The

migration 'Update' feature should 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

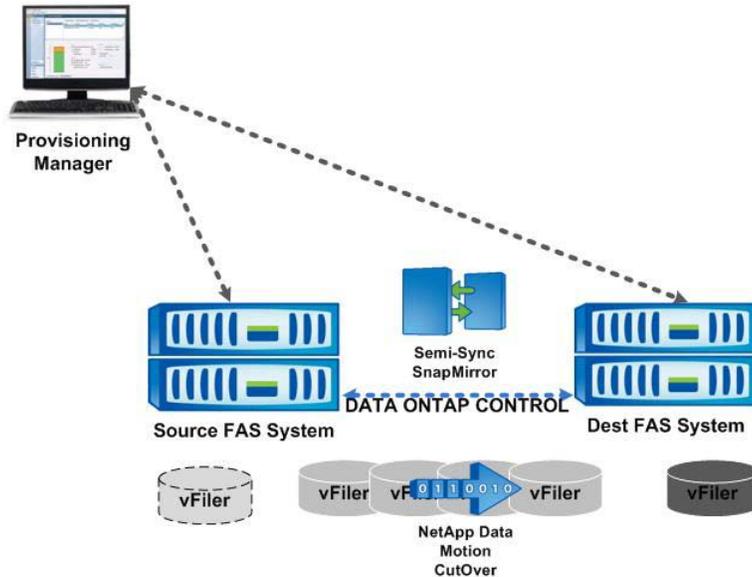


Figure 6) Cutover phase.

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

Data Motion fences some admin operations on the vFiler volumes before cutover. Some typical admin operations that are fenced are: Snapshot creation, file/LUN FlexClone<sup>®</sup> creation, SnapMirror or SnapVault updates on the vFiler volumes and so on. This is required for smooth cutover operation. This fencing continues till cutover completes
3. Cutover process is handled by Data ONTAP:
  - During the cutover, the migration needs to complete within the 120-second window. The two Data ONTAP 7.3.3 systems negotiate with each other to perform the migration. Provisioning Manager is not involved in this part of the process; so its failure during this time will not cause a failure of the Data Motion migration or an application outage.
  - Data ONTAP maintains timers that monitor 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.

During the cutover phase, Provisioning Manager performs several configuration steps (like converting all asynchronous SnapMirror relationships to semi-sync before cutover and waiting for the relationships to synchronize) before starting the actual cutover. These configuration steps could take some time to complete. Effectively, the cutover job in Provisioning Manager may take a long time to complete even though the actual cutover takes maximum of 120 seconds. Fencing of admin operations on the vFiler volumes will continue for whole duration of the cutover process.

**Client experience during online migration:**

- 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 would 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.

**Upon successful completion of NetApp Data Motion cutover:**

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 corresponding to the vFiler volumes must be reconfigured.

**If the NetApp 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:
  - SnapMirror relationships are maintained for cutover retry.
  - SnapMirror relationships are set back to async mode.
  - User can later retry the cutover from Provisioning Manager.
  - 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

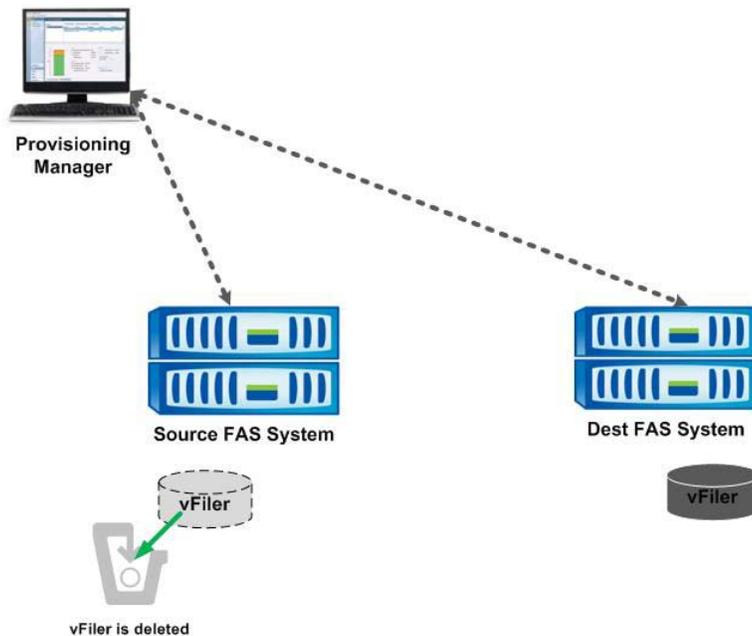


Figure 7) Cleanup phase.

The cleanup phase is executed only when it is initiated from Provisioning Manager.

### During this phase:

Provisioning Manager automatically removes all source FlexVol volumes.

### 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 may be useful if you face any unforeseen issues after cutover (like 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 NetApp 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 NETAPP DATA MOTION WORKFLOW**

- 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 (like cutover exceeding the 120 sec window or the source vFiler unit not coming up within the 120 sec window after an aborted cutover process).

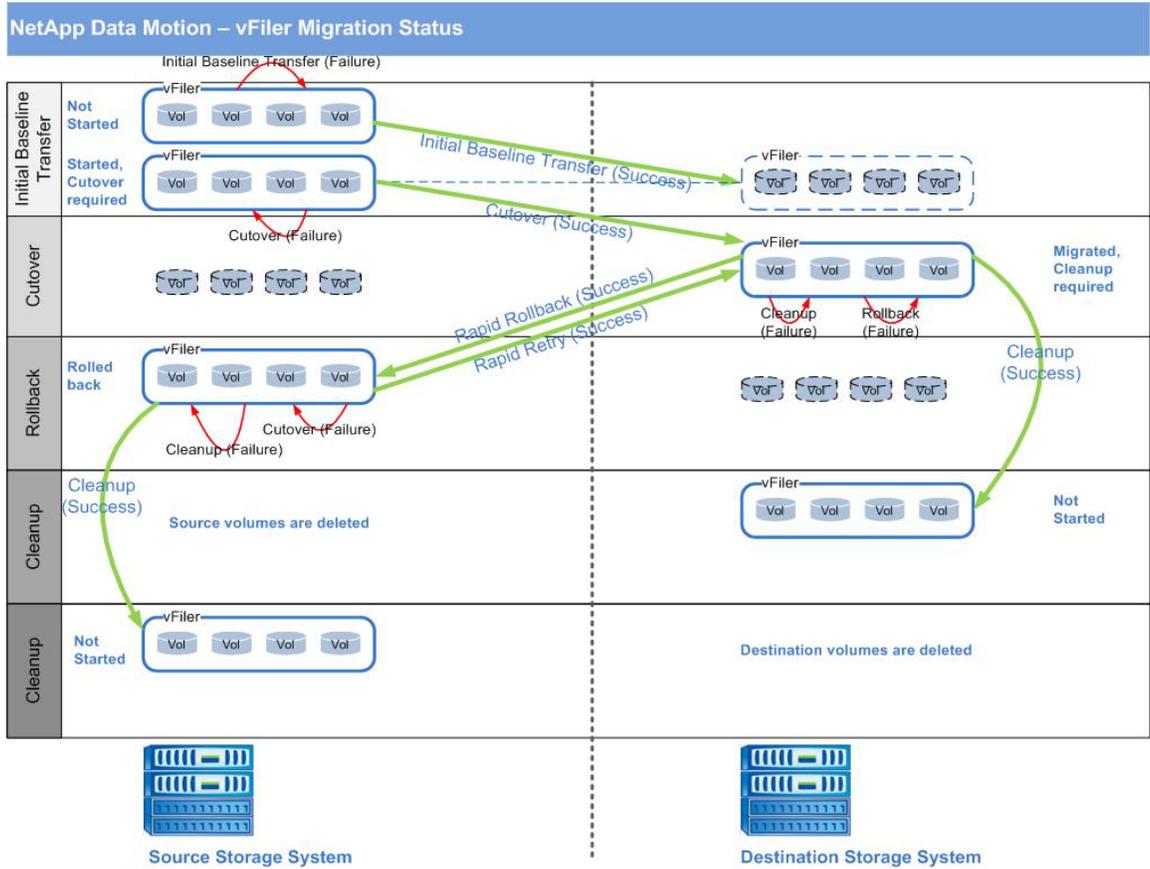


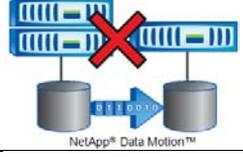
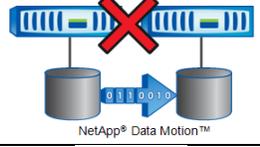
Figure 8) vFiler migration status.

## 7 SUPPORTED CONFIGURATIONS

### 7.1 SUPPORTED MIGRATION SCENARIOS

#### NETAPP DATA MOTION IS NOT SUPPORTED WITHIN A NETAPP CONTROLLER OR A NETAPP HIGH-AVAILABILITY PAIR

Table 3) Supported migration scenario.

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
	Between controller heads of the same high-availability pair	No

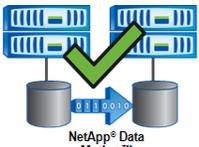
#### STORAGE CONTROLLER RELATED REQUIREMENTS FOR NETAPP DATA MOTION

- Data Motion supports migration of vFiler units from a storage controller head that is part of a NetApp high-availability pair to another controller head that is part of a different NetApp high-availability pair. Both source and destination high-availability clusters should be up and running. If source or destination heads are in takeover/failback mode, then Data Motion is not supported and execution of this scenario is prevented by Provisioning Manager
- Data Motion vFiler migration between the two controller heads of a NetApp high-availability pair is **not** supported, and execution of this scenario is prevented by Provisioning Manager.

- Data Motion migration from a single controller head to another single controller head (which are not part of the any high-availability pair) is **not** supported. Although Provisioning Manager does not prevent the execution of this scenario, NetApp does not recommend this migration scenario.
- Data Motion migration is **not** supported from a single controller head to another controller head that is part of a NetApp high-availability controller pair and vice versa. Although Provisioning Manager does not prevent the execution of these scenarios, NetApp does not recommend these migration scenarios.

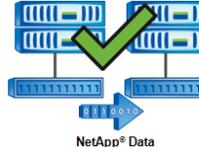
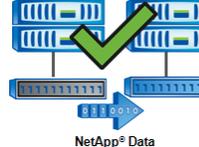
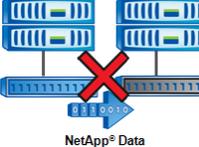
**DATA MOTION IS SUPPORTED ONLY BETWEEN SIMILAR SYSTEM MODELS (THAT IS BETWEEN SYSTEMS HAVING THE SAME AMOUNT OF NVRAM) OR FROM A LOW END SYSTEM MODEL TO A HIGH END SYSTEM MODEL (THAT IS FROM A SYSTEM WITH SMALLER NVRAM TO A SYSTEM WITH LARGER NVRAM)**

Table 4) Supported migration scenario.

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 larger NVRAM to a high-availability controller having smaller NVRAM	No

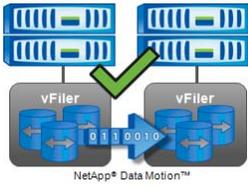
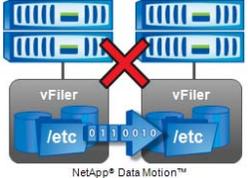
**NETAPP DATA MOTION IS SUPPORTED ONLY BETWEEN SAME SPEED DRIVES OR FROM SLOWER TO FASTER SPEED DRIVES**

Table 5) Supported migration scenario.

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

**NETAPP DATA MOTION IS SUPPORTED ONLY FOR VFILER UNITS THAT OWN FLEXVOL VOLUMES AS STORAGE CONTAINERS**

Table 6) Supported migration scenario.

Migration Scenario	Description	Supported?
 <p>The diagram shows two vFiler units connected by a double-headed arrow. Each vFiler unit contains a FlexVol volume (represented by a blue cylinder) which contains a qtree (represented by a blue folder icon). A green checkmark is placed over the vFiler units, indicating that this configuration is supported for migration.</p>	<p>vFiler units owning FlexVol volumes where an owned FlexVol volume may contain qtrees</p>	Yes
 <p>The diagram shows two vFiler units connected by a double-headed arrow. The left vFiler unit contains a qtree (blue folder icon) that is not inside a FlexVol volume. The right vFiler unit contains a traditional volume (blue cylinder labeled '/etc'). A red 'X' is placed over the vFiler units, indicating that these configurations are not supported for migration.</p>	<ul style="list-style-type: none"> <li>vFiler units owning qtrees whose owning FlexVol volumes reside outside the vFiler unit</li> <li>vFiler units owning traditional volumes</li> </ul>	No

**7.2 SUPPORTED VFILER CONFIGURATIONS**

NetApp Data Motion performs migrations by migrating entire vFiler units, so all FlexVol volumes owned by the vFiler unit must be migrated together in one operation. Table 7 lists the maximum number of FlexVol volumes per vFiler unit supported by NetApp Data Motion corresponding to different storage platforms.

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

Table 7) FlexVol volumes per vFiler unit.

Platform	FlexVol Volumes per vFiler Unit
20X0A	4 FlexVol volumes
30X0A <sup>1</sup>	8 FlexVol volumes
31X0A	8 FlexVol volumes
60X0A	20 FlexVol volumes

<sup>1</sup> NetApp Data Motion supports 4 FlexVol volumes for FAS 3050 platform.

## 7.3 FLEXVOL REQUIREMENTS FOR NETAPP DATA MOTION

### SIZE OF FLEXVOL VOLUMES

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

### SIZE OF AGGREGATES

- To achieve the maximum number of FlexVol volumes per vFiler unit, NetApp recommends that aggregates contain at least 12 disks.

### SNAPLOCK NOT SUPPORTED

- SnapLock volumes cannot be migrated with Data Motion.

### SUPPORT FOR CIFS

- CIFS shares within FlexVol volumes owned by the vFiler unit are migrated as part of Data Motion. However the CIFS sessions are terminated and clients need to 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. Users should be instructed to end their sessions before the cutover.

### SUPPORT FOR DEDUPLICATION AND FILE/LUN FLEXCLONE

- Data Motion is supported with Deduplication or FlexClone at the files or LUN level. However, the combination of having to manage resources associated with non-disruptive migration and metadata for dedupe or FlexClone could result in a small possibility of degradation to the performance of client applications. This may happen only if Data Motion is performed under high load (greater than 60%). Therefore, Data Motion should not be attempted under such conditions.
- NetApp best practice recommendation is that NetApp 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 FlexClone at the files or LUN level (on either the source or destination system). If excessive load is encountered during the migration, NetApp Data Motion can be aborted by the storage administrator, while maintaining access to data.
- For details of the recommended monitoring during the NetApp Data Motion cutover period for systems running dedupe or FlexClone at the files or LUN level, see the following NOW article: <https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb56838>
- 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

- NetApp FlexClone volumes can be members of vFiler units that are migrated with NetApp Data Motion. FlexClone volumes are expanded at the destination, meaning that they will be full-size independent volumes at the destination. Table 8 describes three FlexClone scenarios and their supportability with Data Motion.

Table 8) FlexClone support.

Parent Volume	FlexClone volume	NetApp 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.

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

## 7.4 USING NETAPP DATA MOTION WITH OTHER NETAPP PRODUCTS

Table 9) Data Motion compatibility with other NetApp products.

NetApp Products	Supported for NetApp 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 ' <a href="#">Support for Volume FlexClone</a> '
LUN Cloning	Yes	
Snapshot	Yes	
Fabric MetroCluster	Yes	
Qtree SnapMirror	Yes	Operation is fenced during NetApp Data Motion cutover.
Volume SnapMirror Async	Yes	Operation is fenced during NetApp Data Motion cutover.
SyncMirror	Yes	
SnapManager for Oracle 3.0.3	Yes	Using NFS v3
SnapManager for SAP 3.0.3	Yes	
SnapManager for Virtual Infrastructure (SMVI) 2.0	Yes	
SnapManager for Hyper-V	Yes	

Deduplication	Partial	See section <a href="#">‘Support for Deduplication and File/LUN FlexClone’</a>
File/LUN FlexClone volumes	Partial	See section <a href="#">‘Support for Deduplication and File/LUN FlexClone’</a>
NDMP/Dump	Yes	Operation is fenced during NetApp Data Motion cutover.
PAM I	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 NetApp Data Motion. So after NetApp Data Motion completes, the working set needs to be populated again on the destination storage system (i.e. PAM cache will need to be "warmed" back up) and throughput may degrade during this time.
PAM II	Yes	
Volume SnapMirror Sync	Partial	Only in default vFiler interfaces. Also operation is fenced during NetApp Data Motion cutover.
SnapVault	Partial	A SnapVault source can be a vFiler unit, but the destination should not.
FlexCache	No	Only FlexCache source volumes can be migrated with Data Motion.
SnapLock Compliance	No	SnapLock volumes are not supported for NetApp Data Motion
SnapLock Enterprise	No	SnapLock volumes are not supported for NetApp Data Motion
SnapManager 5.0 for Microsoft Exchange	Yes	
SnapManager 5.0R1 Microsoft SQL Server	Yes	
SnapManager 2.0 for Microsoft Office SharePoint Server	Yes	
Thin Provisioning	Yes	

## 8 IP NETWORKING REQUIREMENTS FOR NETAPP DATA MOTION

With the vFiler unit being the unit of migration in Data Motion, all MultiStore specific networking requirements must be met. For details about MultiStore networking methods, see the *Data ONTAP 7.3.3 MultiStore Management Guide* available on [NOW](#).

There are several ways to assign network addresses to vFiler units. A vFiler unit has interfaces assigned to it as part of the provisioning process. These interfaces can be any of the following: physical interfaces, IP aliases, VLANs, or virtual interfaces (VIFs) - either using the default or a non-default IP space. NetApp does not recommend using physical interfaces, because the system will quickly run out of available interfaces. Table 10 illustrates the different 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.

Table 10) Data Motion compatibility with other NetApp products.

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

### USING IP ALIASES

For assigning IP addresses to the vFiler units, multiple IP aliases are created on a network interface of the physical FAS system and each vFiler unit is assigned an IP alias. All clients of the vFiler unit must make the request on this IP alias. NetApp recommends this method for situations where vFiler units do not need physical or logical network separation.

### USING VLANS AND VIFS

Both types of virtual interfaces 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. See [Appendix B](#) for more information.

### 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.
- Each network interface (physical, VIF, or VLAN interface) of the FAS system can belong only to one IPspace; however one IPspace can have multiple network interfaces.
- IPspaces are created when vFiler units need to have their own secure storage, administration, and routing leading to secure multi-tenancy.
- IP addresses 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. Also each 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 the remaining 100 can be created by the user. 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 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 setup.
- VLAN tagging with IPspaces make sure that network packets are forwarded to the appropriate IPspace and securely delivered to the vFiler unit contained inside the IPspace.

For Data Motion, using VLAN tagging with IPspaces is the recommended method.

For a step-by-step procedure, see [Appendix B](#). Appendix B also covers network configuration best practices for network switches, storage systems, and vFiler units.

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

## 8.1 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](#).

1. Configure VIF on the storage controller. In this example, VIF named `smvif36` is used.
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. Now assign the interface to the IPspace. In the following example, the single mode VIF is assigned to the IPspace 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 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 the 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.

## 8.2 NETWORK CONFIGURATION SCENARIOS FOR NETAPP DATA MOTION

### VFILER UNITS USING IP ALIASES

1. Consider a typical IP storage infrastructure supporting Data Motion. Each NetApp array 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.

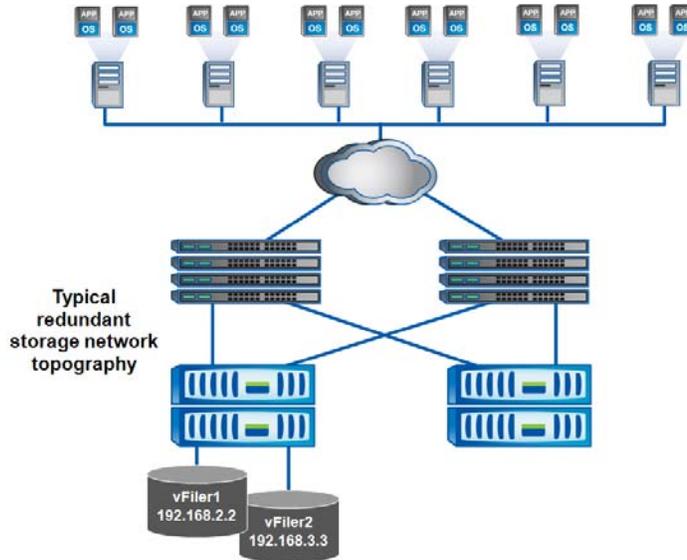


Figure 9) Sample Topography of an IP storage infrastructure.

2. In this example, IP aliases are assigned to the vFiler units on the NetApp array. Each one is on a different 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 above to access the vFiler units.

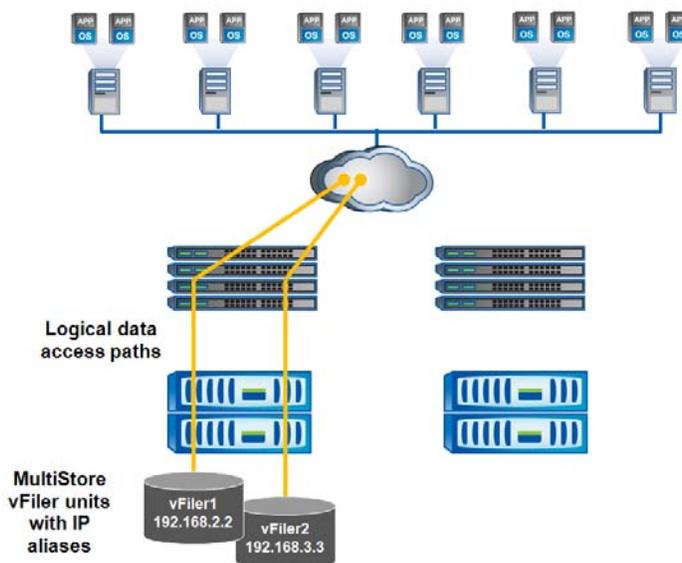


Figure 10) Network configuration for Data motion.

- 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. The clients are able to access data in the new location. From the network's perspective movement of the IP address is very similar to how an IP address would move when a link in a single-mode VIF fails and the IP address is moved to another port on the network.

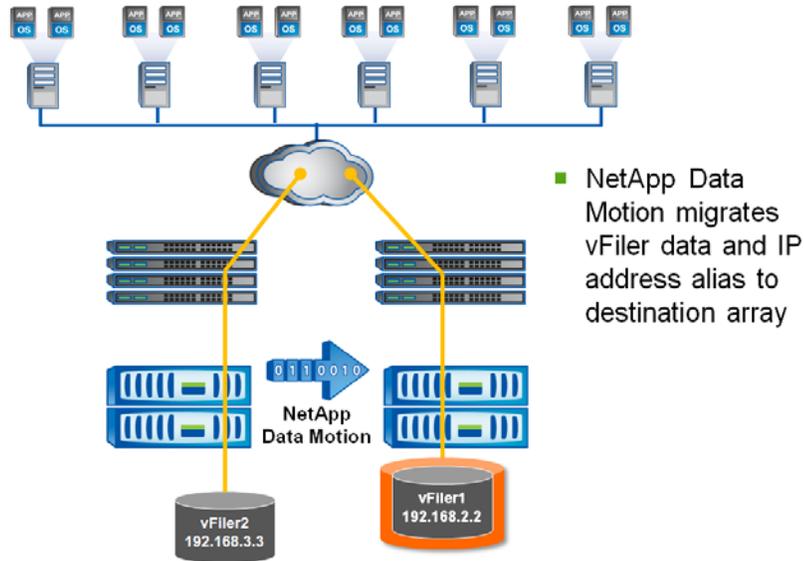


Figure 11) Network configuration for Data motion.

## VFILER UNITS USING VLAN TAGGING AND IPSPACES

1. 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 may be sharing the same physical ports on 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 shown in the following figure is the network configuration for Data Motion recommended by NetApp. Refer to [Appendix B](#) for an end-to-end process of configuring the network for Data Motion. Appendix B covers network configuration best practices for network switches, storage systems, and vFiler units.

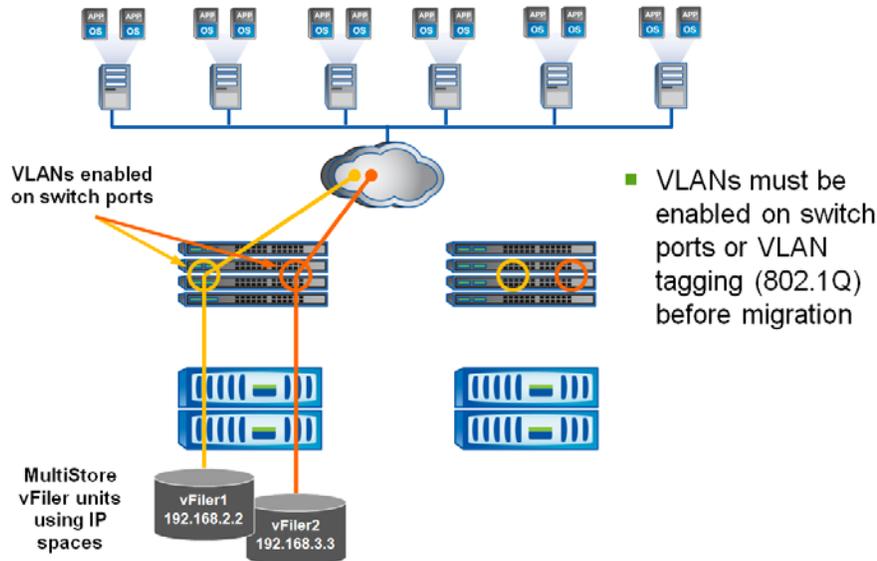


Figure 12) vFiler units using VLAN tagging and IPSpaces.

### Best Practices:

- Both 802.1q (VLAN tagging) and 802.3ad (Trunking) standards should be used together to enable VLANs to move without disruption.
  - When using IPSpaces and VLAN tagging, special attention needs to be taken to make sure that the team managing the switches have 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 will never be accessed by that VLAN prior to the Data Motion migration.
2. When migration is performed, Data Motion creates VLANs on the destination NetApp array, but it cannot configure the VLAN trunking of ports on the network switch. Therefore, it is very important to configure this before migration. Data Motion is not capable of checking the switch to make sure that the VLANs are properly configured. If the VLANs are not configured properly on the switch connected to the destination array, the Data Motion migration will complete, however, the clients will not be able to access the vFiler unit at the new location, and this could cause an application outage.

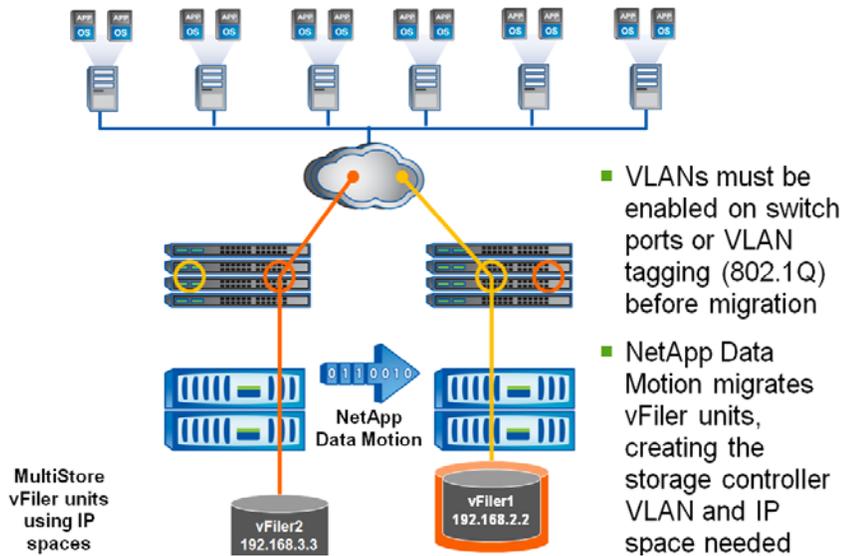


Figure 13) vFiler units using VLAN tagging and IPSpaces.

### 8.3 OTHER NETWORKING CONSIDERATIONS

- Before migration, make sure that the destination NetApp system is accessible by clients of the migrating vFiler unit.
  - VLANs must be preenabled 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 the FlexVol volumes assigned to the vFiler unit and NetApp Data Motion will intelligently distribute this bandwidth across the vFiler volumes.
- Data Motion replicates the vFiler unit network configuration from source to destination; this 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

## 9 NETAPP DATA MOTION BEST PRACTICES

### 9.1 COMPONENT-LEVEL BEST PRACTICES: MULTISTORE, SNAPMIRROR, PROVISIONING MANAGER

- NetApp recommends not making any configuration changes until the NetApp Data Motion cutover phase completes. If any configuration change related to the vFiler unit has been done after starting the initial base line transfer, perform a Data Motion update before proceeding to the cutover phase.
- As the entire Data Motion process is controlled through Provisioning Manager, do not make any manual modification to any configuration related to the vFiler units or SnapMirror (for example, the `snapmirror.conf` file).
- Before starting the cutover process from Provisioning Manager, it is recommended to refresh the source and the destination storage systems in the Operations Manager by using the `dfm host discover` command.
- Refer to the SnapMirror Asynchronous and Semi-Synchronous Best Practices:  
<http://www.netapp.com/us/library/technical-reports/tr-3446.html>  
<http://www.netapp.com/us/library/technical-reports/tr-3326.html>

### 9.2 PROCESS-LEVEL BEST PRACTICES

- If there is a long gap between initial baseline transfer phase completion and the subsequent cutover initiation, NetApp recommends first performing an 'Update' before proceeding to the cutover.
- As a best practice, retain a stable backup of all storage entities (volumes and LUNs) 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 NetApp 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 type of operations after successful migration, NetApp recommends carefully checking 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 having the same speed disk drives.
- The current release of NetApp 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, NetApp recommends carefully checking 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.

### 9.3 PERFORMANCE CONSIDERATION

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.

NetApp recommends that during Data Motion cutover the CPU and disk I/O load on the source and destination storage systems should be in low to medium range. Since semi-sync SnapMirror is involved in the cutover phase, additional headroom should be kept for the CPU and disk I/O load before initiating cutover.

**BEST PRACTICE:** While choosing a Data Motion destination, make sure that it'll be able to handle the workload being migrated.

In this initial release of NetApp 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, NetApp Data Motion can be aborted by the storage administrator. Refer section '[Support for Deduplication and File/LUN FlexClone](#)' for additional details and recommended monitoring methods during migration where Dedupe or FlexClones® at the Files or LUN level exist on either the source or destination system.

#### 9.4 IP NETWORK CONFIGURATION BEST PRACTICES

- NetApp recommends using VLAN trunking on the network switches of any storage system that may be a source or destination for Data Motion.
- All required VLANs must be preenabled on destination switch ports. Before migration, make sure that the destination FAS system is accessible by clients of the migrating vFiler unit.
  - When the migration is performed, Data Motion creates the VLANs on the destination NetApp array, but it cannot configure the VLAN trunking of ports on the network switch. Therefore it is very important to configure this before migration. Data Motion is not capable of checking the switch to make sure that the VLANs are properly configured. If the VLANs are not configured properly on the switch 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 could cause an application outage.
- As a best practice, the network of the vFiler unit should be configured inside an IP space associated with a VLAN interface over a VIF created on the storage system.
- NetApp recommends using a separate IPspace per vFiler unit (corresponding to different tenants, departments, or applications).
- For details, see "[IP Networking Requirements for NetApp Data Motion](#)" and "[Other Networking Considerations](#)," earlier in this document.
- [Appendix B](#) describes a step-by-step networking best practices implementation.

#### 9.5 STORAGE-LEVEL BEST PRACTICES

##### REQUIRED DATA ONTAP CONFIGURATION

The following Data ONTAP licenses must be installed for NetApp Data Motion: iSCSI and/or NFS, MultiStore, SnapMirror, Synchronous SnapMirror.

For a complete list of hardware and software requirements, see [Supported Platforms and Requirements](#).

##### SUPPORT FOR DEDUPLICATION AND FILE/LUN FLEXCLONE

NetApp Data Motion supports deduplication and FlexClone at the Files or LUN level under low to medium load. For fastest migration times and minimal impact, NetApp recommends performing Data Motion during non peak times, or periods during lower load. For details and recommended monitoring during migration when deduplication or FlexClone at the files or LUN level exist on either the source or destination system, see section [Support for Deduplication and File/LUN FlexClone](#).

##### VOLUME FLEXCLONE SUPPORT

FlexClone volumes can be members of vFiler units that are migrated with Data Motion. For details, see section [Support for Volume FlexClone](#).

##### DISK TYPE SUPPORT—FC, SATA, SAS

NetApp Data Motion supports migration between disks with the same speed or from slower to faster disks. Table 11 shows the supported configurations.

Table 11) Data Motion compatibility with other NetApp products.

Source Aggregate disk type	Destination Aggregate disk type	NetApp 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
ATA	FC	Allowed
BSAS	FC	Allowed
SATA	SAS	Allowed
ATA	SAS	Allowed
BSAS	SAS	Allowed

#### SNAPDRIVE AND SNAPMANAGER BEST PRACTICE

- Do not create or delete any LUN inside volumes associated with the vFiler unit after starting the migration process and before cutover completes. NetApp recommends performing all vFiler resource and dataset related change operations before the initial baseline transfer.
- During the cutover or rollback, any backup and restore operations would be unsuccessful for SnapManager products (SME, SMSQL, SMVI, SMHV, SMO, SMSAP, and SMOSS) as there are no retry options from SnapDrive (SDW and SDU). It is recommended to refrain from using any backup and restore command during these phase (cutover or rollback).

#### HOST OPERATING SYSTEM DISK TIME OUT SETTINGS

- Set the disk timeouts in the host operating systems by using the NetApp Host Utilities Kit (HUK). Refer to the NetApp Host Utilities *Installation and Setup Guide* for your host operating system (Windows<sup>®</sup>, ESX, Linux<sup>®</sup>, and so on) for the procedure.
- Sample disk timeout values: 190 sec for a guest operating system running on VMware ESX, 120 sec for a standalone Windows host without MPIO, and 20 sec or 60 sec (for HUK v5.1 and after) with MPIO.

#### CFO

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

#### METROCLUSTER

- NetApp Data Motion does not support cutover when one storage node of the MetroCluster system has taken over the partner either automatically or manually (CFOD). If the takeover event happens after the cutover operation has started, the cutover will fail.

## 9.6 BEST PRACTICES FOR SUPPORTING SERVER VIRTUALIZATION ENVIRONMENT

### VMWARE

- As mentioned earlier, set the disk timeout values in the guest operating systems running on a VMware ESX server by using the NetApp Host Utilities kit.
- If Provisioning Manager is used to create and assign storage volumes (NFS) to vFiler units, those NFS volumes are exported as qtrees. When VMware data stores are created corresponding to these NFS exports, due to a limitation in Provisioning Manager, the storage savings obtained through various NetApp 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 created by Provisioning Manager 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. NetApp best practice recommendation for VMware VI3 and vSphere is to use FlexVol volumes instead of qtrees for VMware NFS data stores. This leads to simplified architecture and also makes the value of NetApp's storage efficiency technologies visible to the VMware server administrator. See [TR3428](#) and [TR3749](#) for details.

## APPENDIX A: THE NETAPP DATA MOTION PROCESS

### OBJECTIVE OF THIS EXERCISE

Use NetApp Data Motion to perform an online migration of a vFiler unit that is serving live host and application data.

### TASKS

Table 12) Steps to perform online migration.

Step	Procedure
1	Create a vFiler unit and provision storage: Method 1-Using NetApp storage system CLI. Method 2-Using NetApp FilerView interface. Method 3-Using NetApp 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.

### CREATE VFILER UNIT AND PROVISION STORAGE

There are three ways in which the vFiler unit can be created and assigned storage. Use one of these methods according to the environment and requirements.

#### METHOD 1: USING THE NETAPP STORAGE SYSTEM CLI

The NetApp storage system CLI offers a rich set of commands for creating and configuring vFiler units. Users 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

Table 13) vFiler command set.

vFiler Command Set
<code>vfiler help</code> - Help for vfiler command
<code>vfiler context</code> - Set the vfiler context of the CLI
<code>vfiler create</code> - Create a new vfiler
<code>vfiler rename</code> - Rename an existing vfiler
<code>vfiler destroy</code> - Release vfiler resources
<code>vfiler dr</code> - Configure a vfiler for disaster recovery
<code>vfiler add</code> - Add resources to a vfiler
<code>vfiler remove</code> - Remove resources from a vfiler

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

For details on the procedure to create vFiler and provision storage using NetApp storage system CLI, see the *MultiStore Management Guide* corresponding to your version of Data ONTAP.

### METHOD2: USING THE NETAPP FILERVIEW INTERFACE

The NetApp FilerView<sup>®</sup> graphical user interface supports the creation and configuration of vFiler units. Users 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

For details on the procedure to use the NetApp FilerView GUI, and to use it to create vFiler units and provision storage, see the *System Administration Guide* and *MultiStore Management Guide* for your version of Data ONTAP.

### METHOD 3: USING NETAPP PROVISIONING MANAGER

NetApp Provisioning Manager automates the process of vFiler creation and storage provisioning. See the *Provisioning Manager 4.0 Workflow Guide* and *Provisioning Manager and Protection Manager Printable Help* for a detailed description of vFiler management using Provisioning Manager.

Figure 14 depicts the high-level steps in the automated creation and provisioning of vFiler units by using Provisioning Manager.

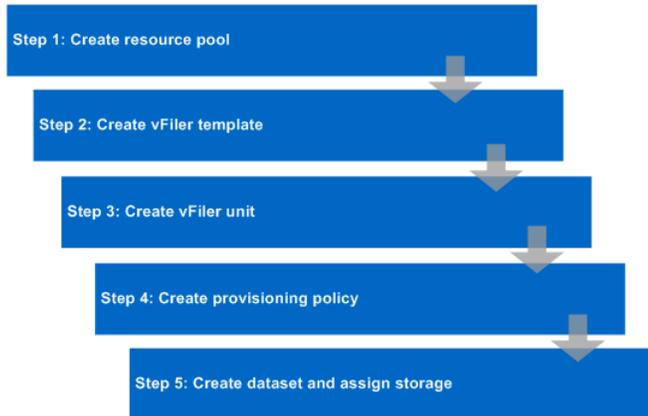


Figure 14) Steps to create and provision vFiler units.

## ASSIGN THE STORAGE FROM THE VFILER UNIT TO THE HOSTS AND APPLICATIONS

Once the vFiler unit has been created and configured, the storage (volumes and LUNs) assigned to the vFiler unit can be mapped to the hosts and the applications, similarly to any NetApp storage system. For details, see the *MultiStore Management Guide*, *File Access and Protocols Management Guide*, and *Block Access Management Guide* for your version of Data ONTAP.

## VFILER UNIT ONLINE MIGRATION USING PROVISIONING MANAGER

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

### INITIAL BASELINE TRANSFER

1. Select the vFiler unit and click 'Start migration' from 'Hosts → vFiler Units'

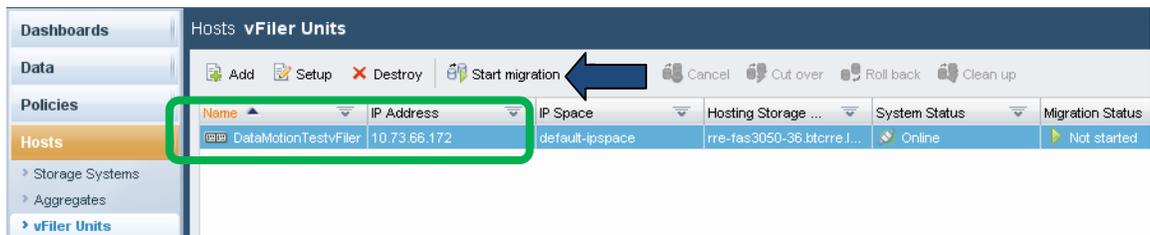


Figure 15) Initial baseline transfer (step1).

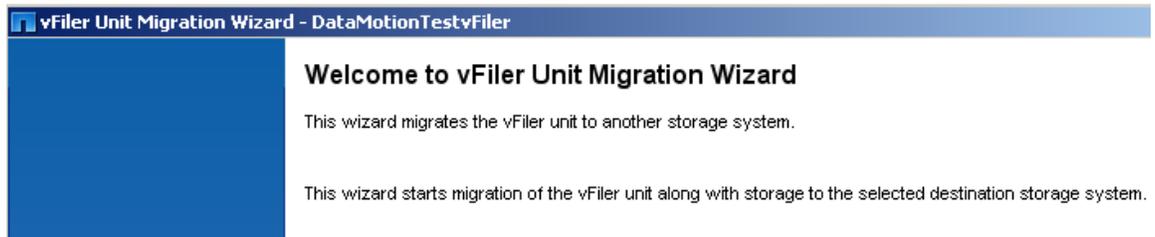


Figure 16) vFiler migration wizard welcome screen.

2. Select Online Migration.

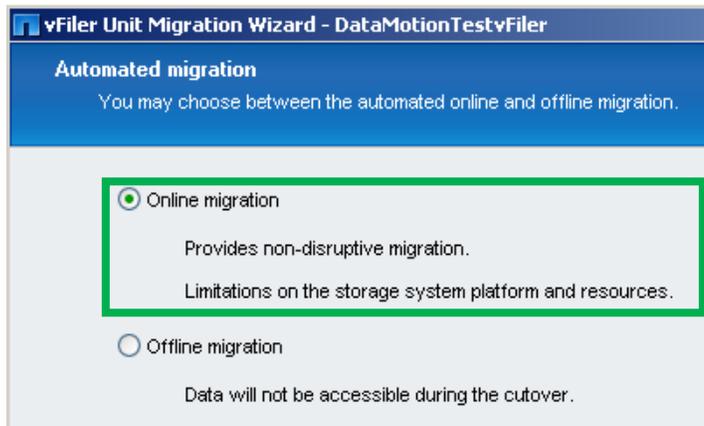


Figure 17) Initial baseline transfer (step 2).

3. Select the destination FAS system to which the vFiler unit is to be migrated.

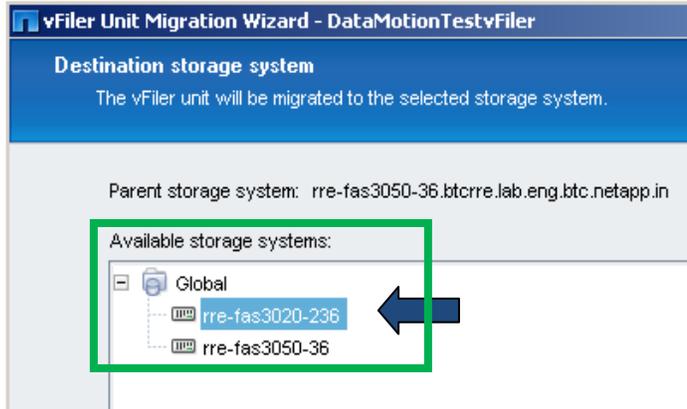


Figure 18) Initial baseline transfer (step 3).

4. The network interface on the destination FAS system should be configured carefully for online migration. Select the network interface to use and the VLAN ID, if any.

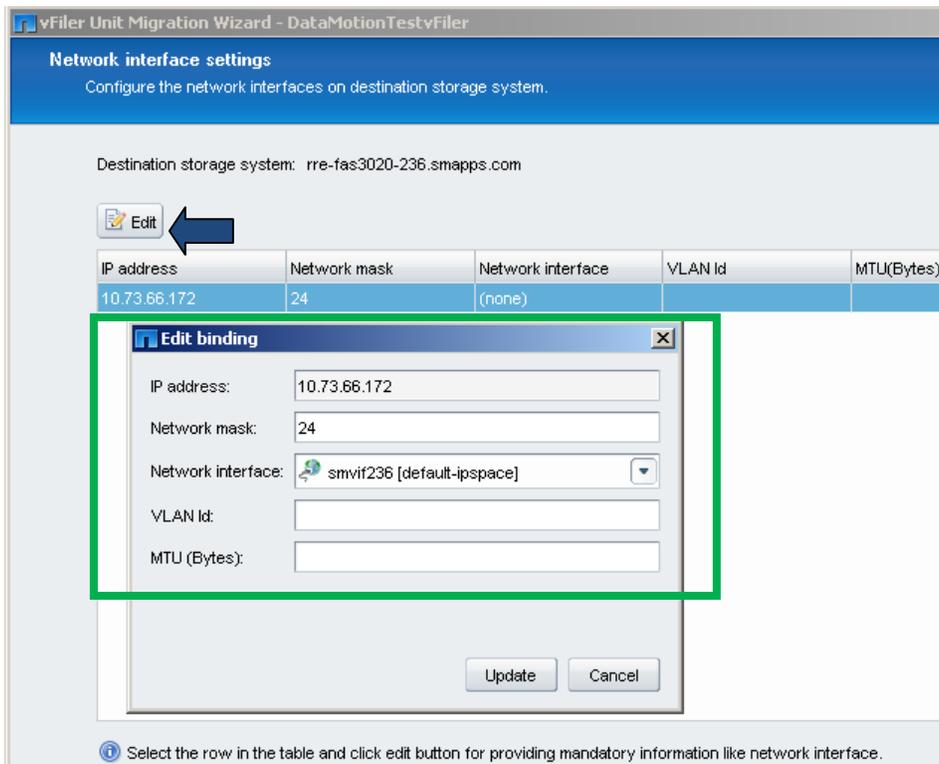


Figure 19) Initial baseline transfer (step 4).

- (Optional) Throttle the bandwidth to be used for the initial baseline transfer.



Figure 20) Initial baseline transfer (step 5).

- There is also an optional feature to cutover to the destination FAS system automatically after initial data transfer. This feature provides the option of 'One Step Migration'. For this example, this option is left unchecked.

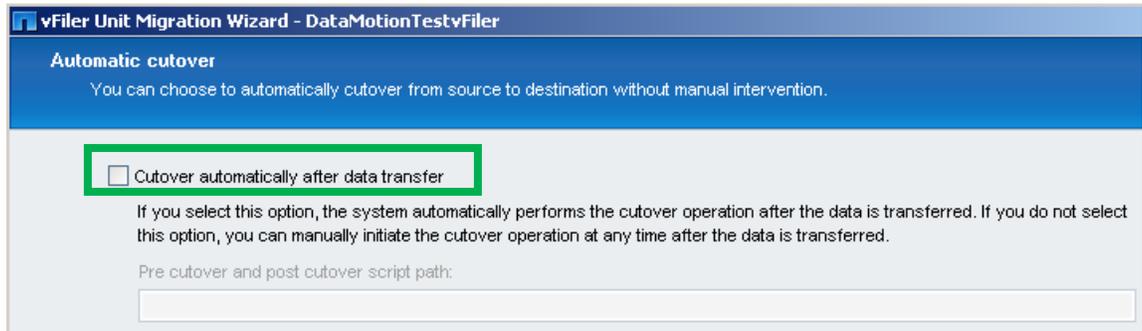


Figure 21) Initial baseline transfer (step 6).

7. Preview the steps that will be automated by Provisioning Manager for the vFiler migration process.

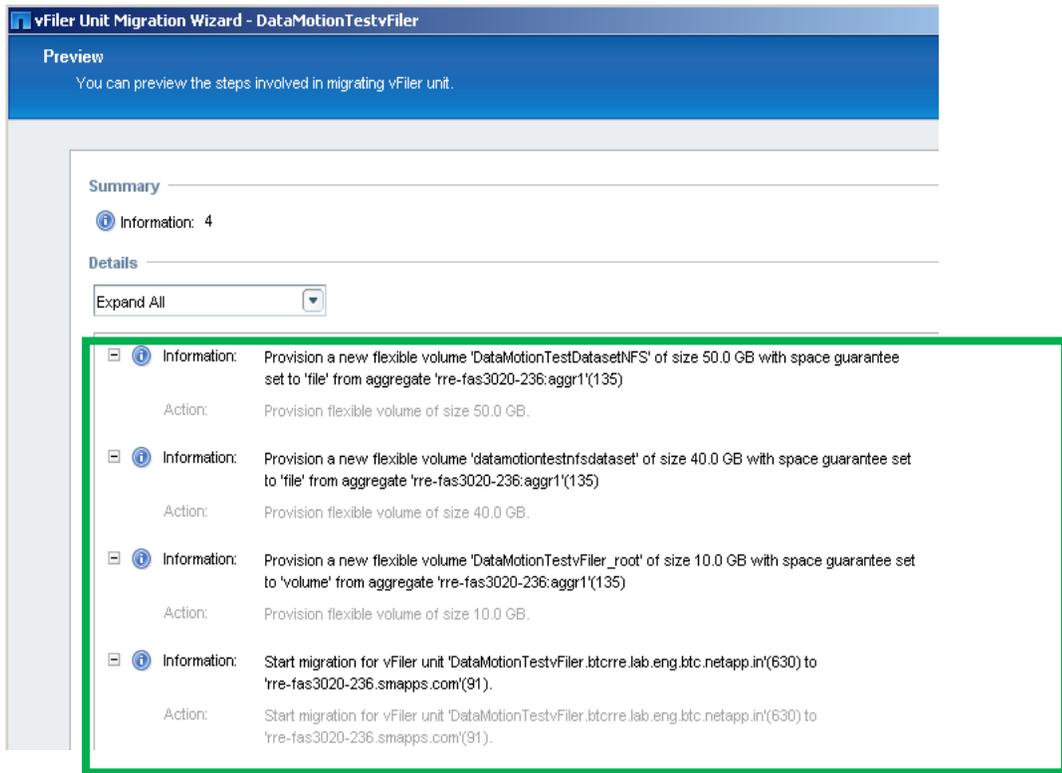


Figure 22) Initial baseline transfer (step 7).

8. Complete the vFiler Unit Migration Wizard.



Figure 23) Initial baseline transfer (step 8).

A monitor window opens, displaying the status of the migration job.

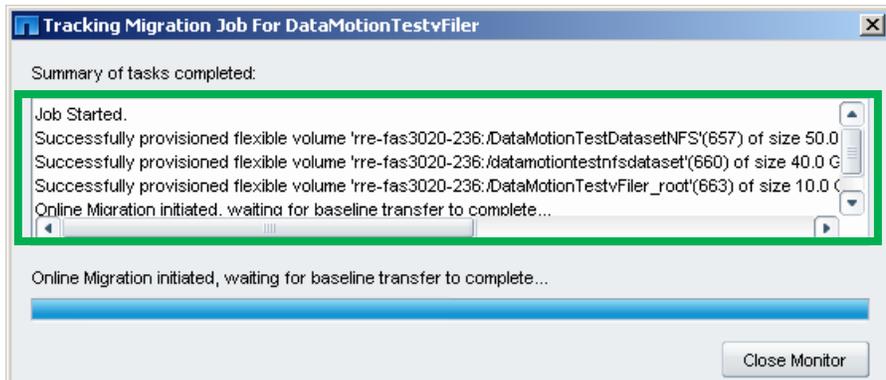


Figure 24) Initial baseline transfer (step 9).

9. The status of the migration job can also be viewed from Data →Jobs.

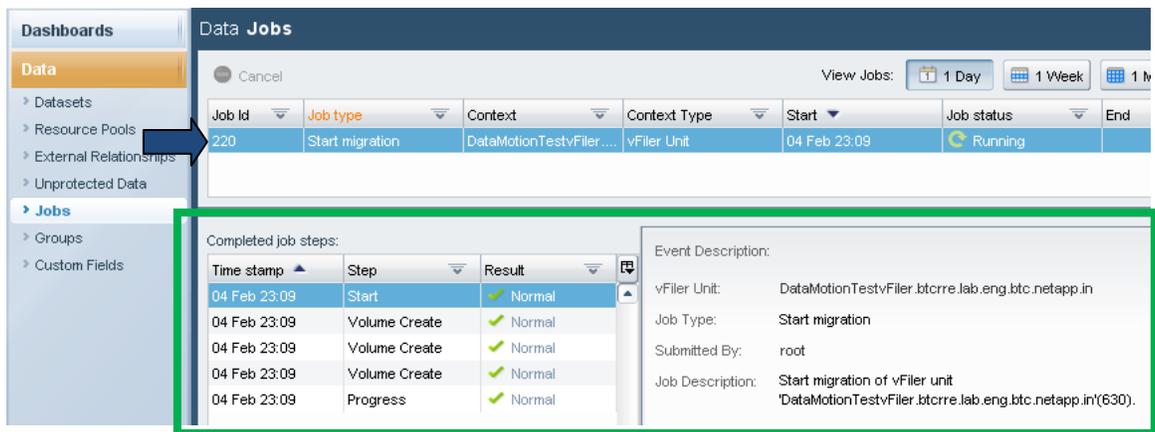


Figure 25) Initial baseline transfer (step 10).

When the baseline transfer of the data is complete, the migration status of the vFiler unit changes to Started, Cutover Required.

At this point, the user has three options:

- 'Update' the baseline created in the destination system (recommended if there is a long gap between the time when baseline transfer was completed and the time when the cutover process is started).
- Cancel the Data Motion process.
- Start the cutover process to migrate the vFiler unit to the destination system.

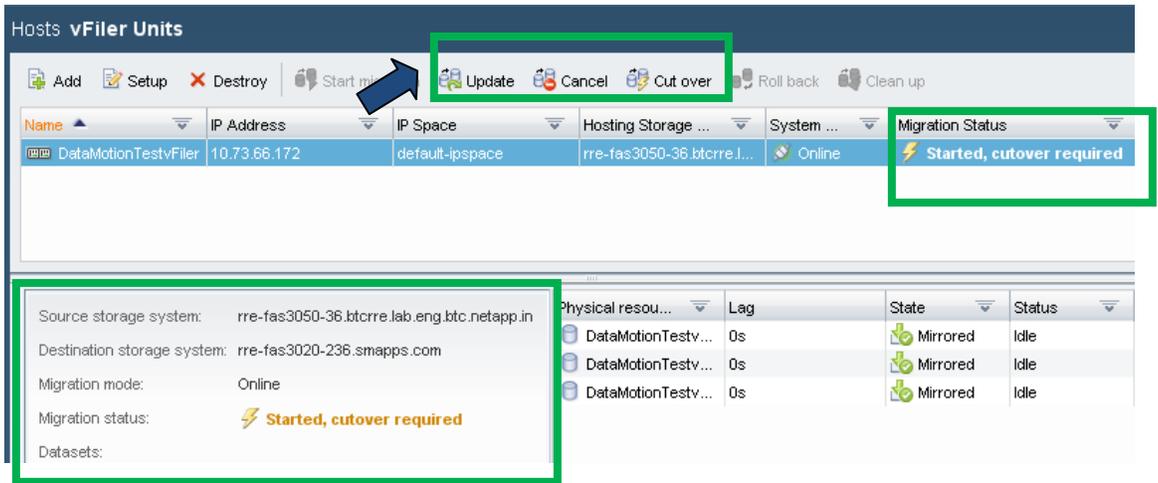


Figure 26) Initial baseline transfer (step 11).

## CUTOVER

1. To start the cutover process, select the vFiler unit from Hosts → vFiler Units and click Cut over.



Figure 27) Migration cutover wizard (1).

2. (Optional) Provide the path of the script to run before and after cutover.
3. Select an option for 'Network Routes Migration.'

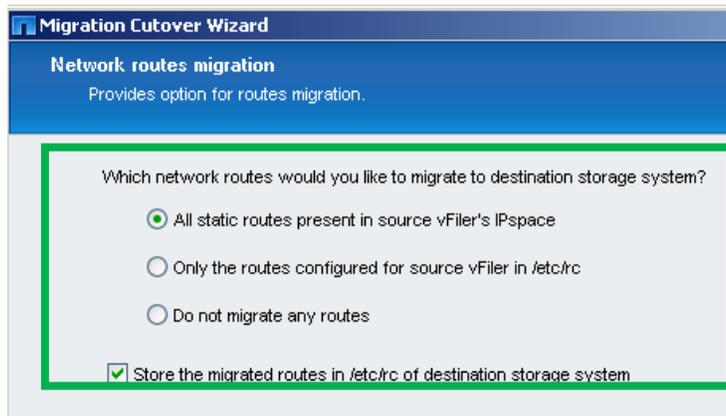


Figure 28) Migration cutover wizard (2).

4. Review the summary of steps that will be automated by Provisioning Manager for cutover to happen.

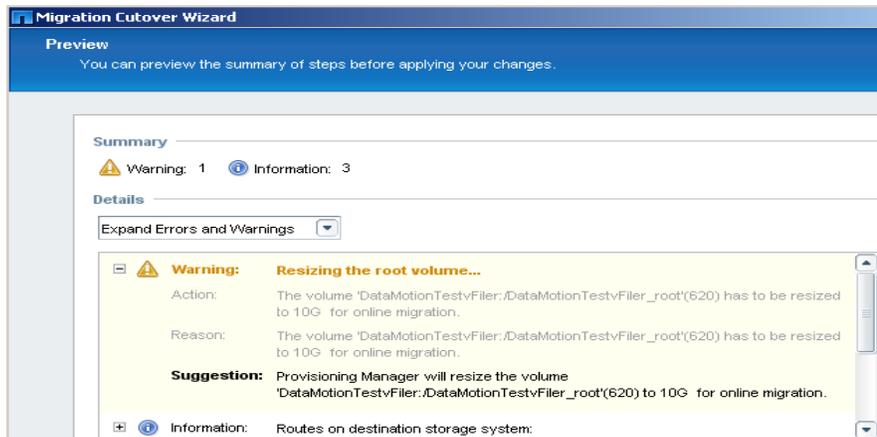


Figure 29) Migration cutover wizard (3).

5. Complete the Migration Cutover Wizard.

A monitoring window opens, displaying the status of the cutover job.

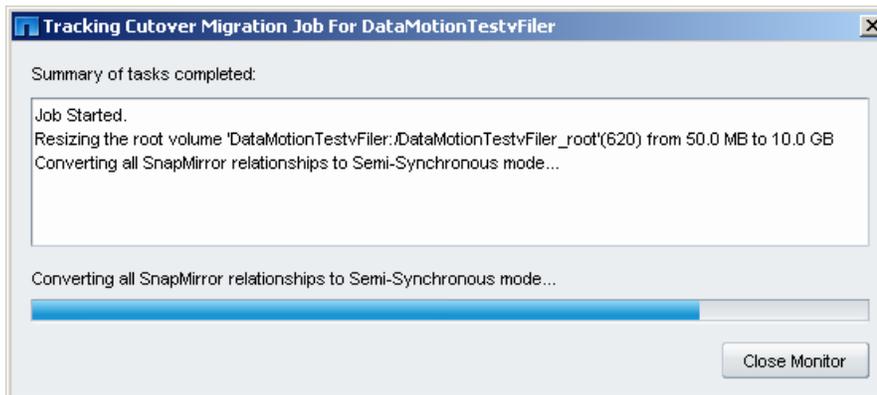


Figure 30) Migration cutover wizard (4).

6. The status of the cutover job can also be monitored from Data → Jobs → Job Type: Cutover.

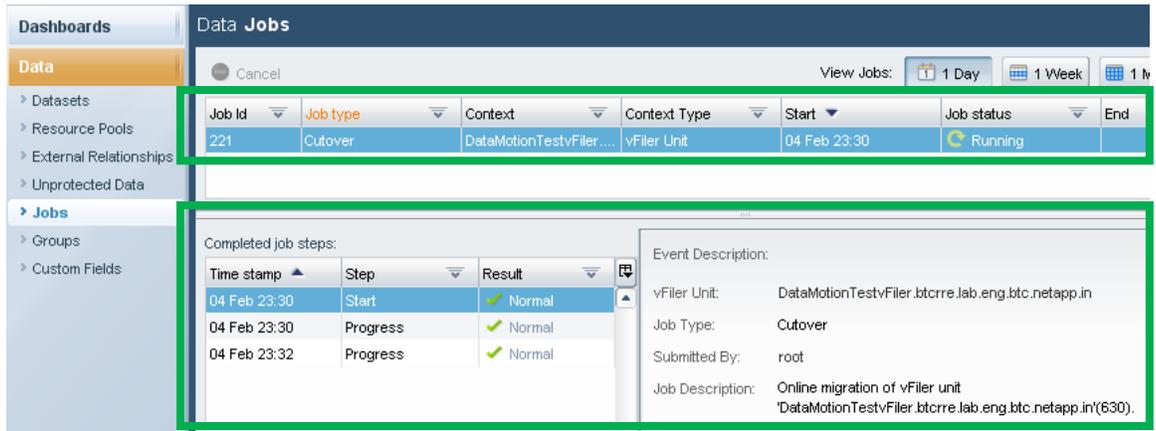


Figure 31) Migration cutover wizard (5).

When the cutover has successfully completed, the migration status of the vFiler unit changes to “Migrated, Cleanup Required”. At this point, the user has the option to either Rollback the vFiler unit to the source storage system or Cleanup the vFiler unit resources from the source storage system.

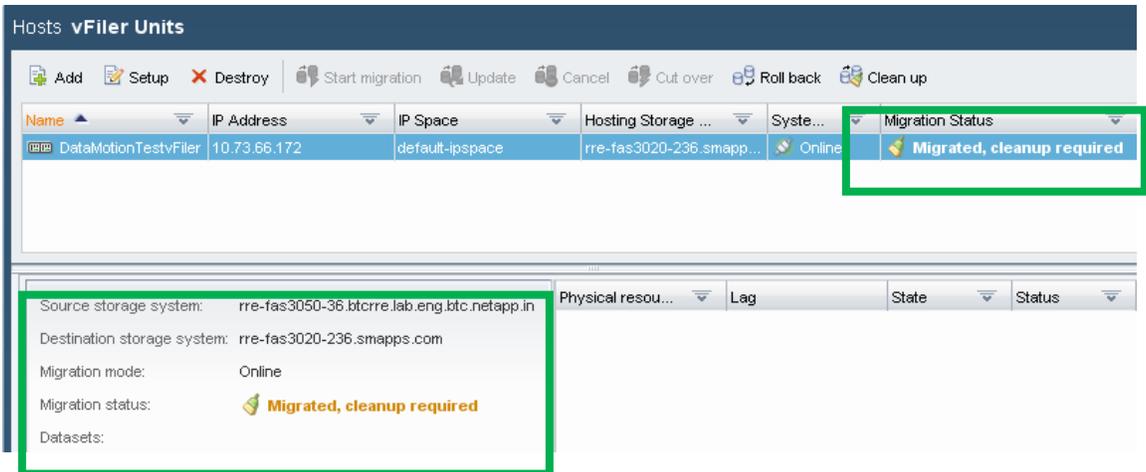


Figure 32) Migration cutover wizard (6).

## ROLLBACK

1. If the user chooses to Rollback the vFiler unit to the source storage system, use Rollback option in the Hosts → vFiler Units.



Figure 33) Migration rollback wizard (1).

- (Optional) Provide a path of the script that needs to be run before and after rollback.
- Select the required options for Network routes migration.

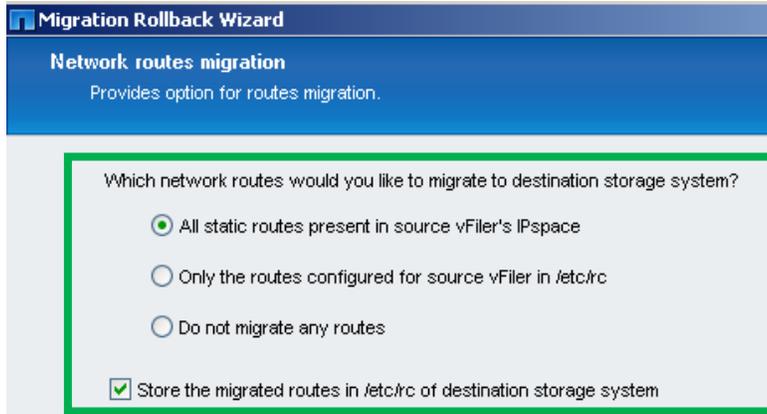


Figure 34) Migration rollback wizard (2).

- Review the steps that will be automated by the Provisioning Manager for Rollback.

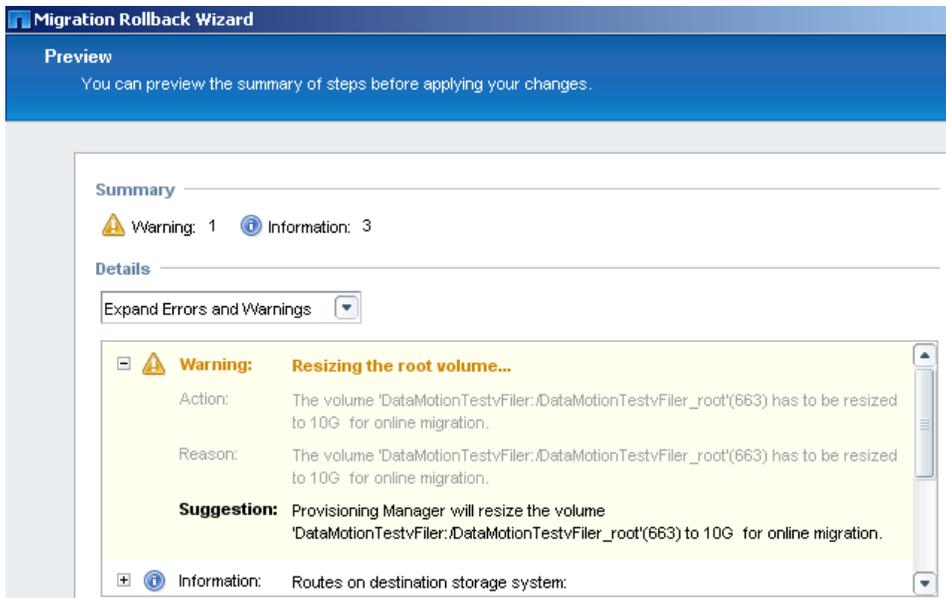


Figure 35) Migration rollback wizard (3).

- Complete the Migration Rollback Wizard.  
A monitoring window opens, displaying the status of the rollback job.

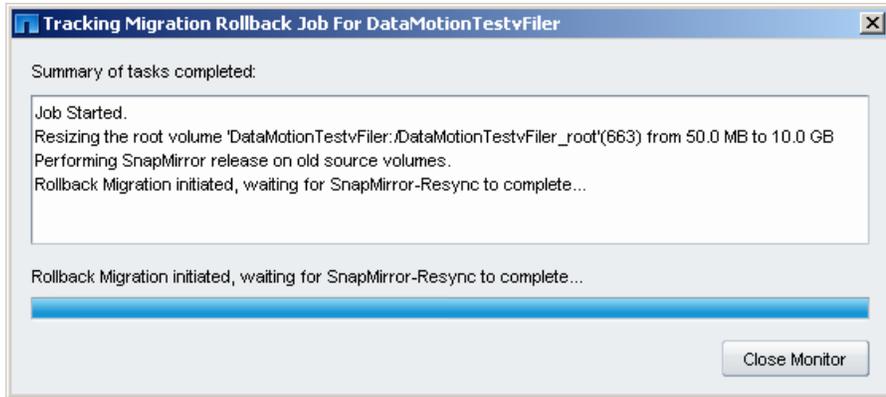


Figure 36) Migration rollback wizard (4).

- The status of the Rollback job can also be monitored from: 'Data → Jobs → Job Type: Migration rollback.

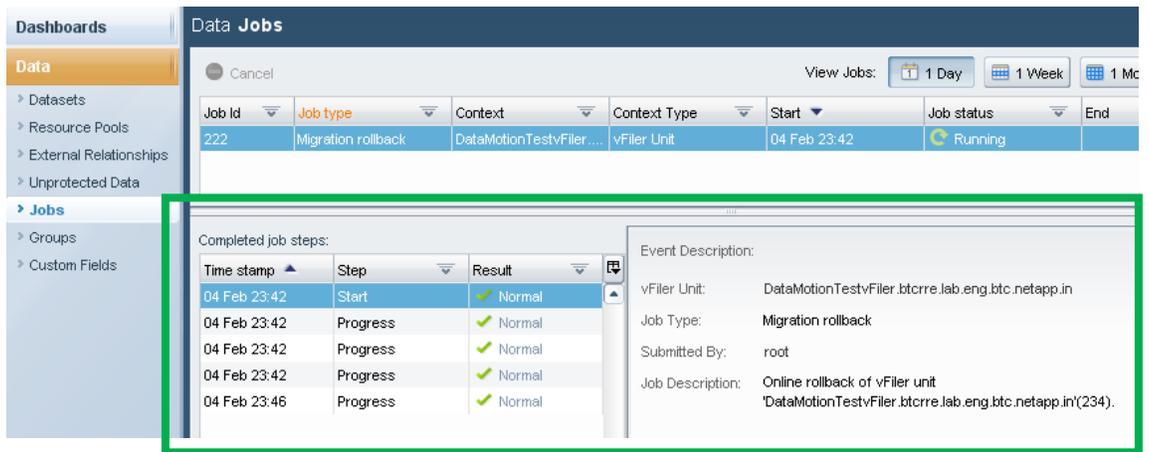


Figure 37) Migration rollback wizard (5).

When rollback is completed successfully, the Migration Status of the vFile unit changes to Rolled back. At this point, the user has the option to either 'cutover' the vFile unit to the destination storage system or to clean up the vFile unit resources from the destination storage system.

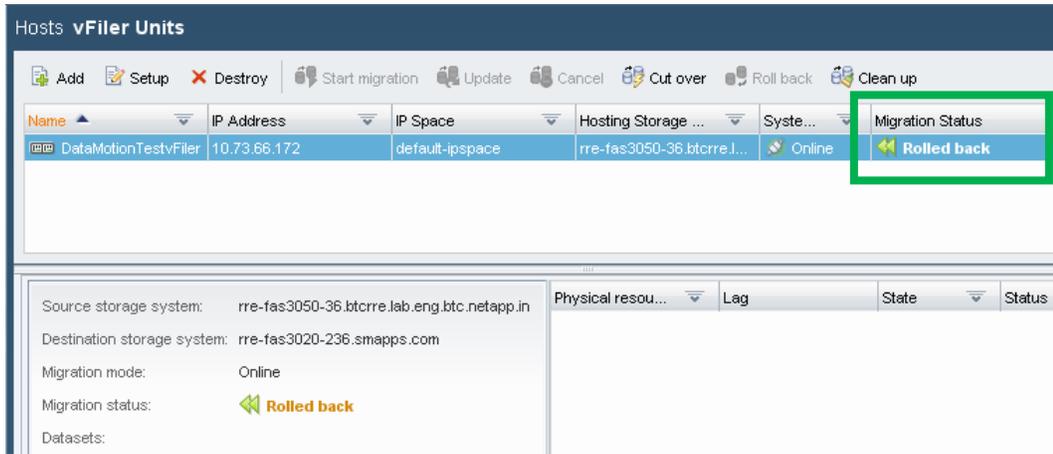


Figure 38) Migration rollback wizard (6).

## CLEANUP

1. After a successful cutover to the destination storage system (step) or after a successful rollback to the source storage system (step), the vFile unit and its resources from the source or destination storage system can be cleaned up using Clean up option in Hosts → vFile Units.

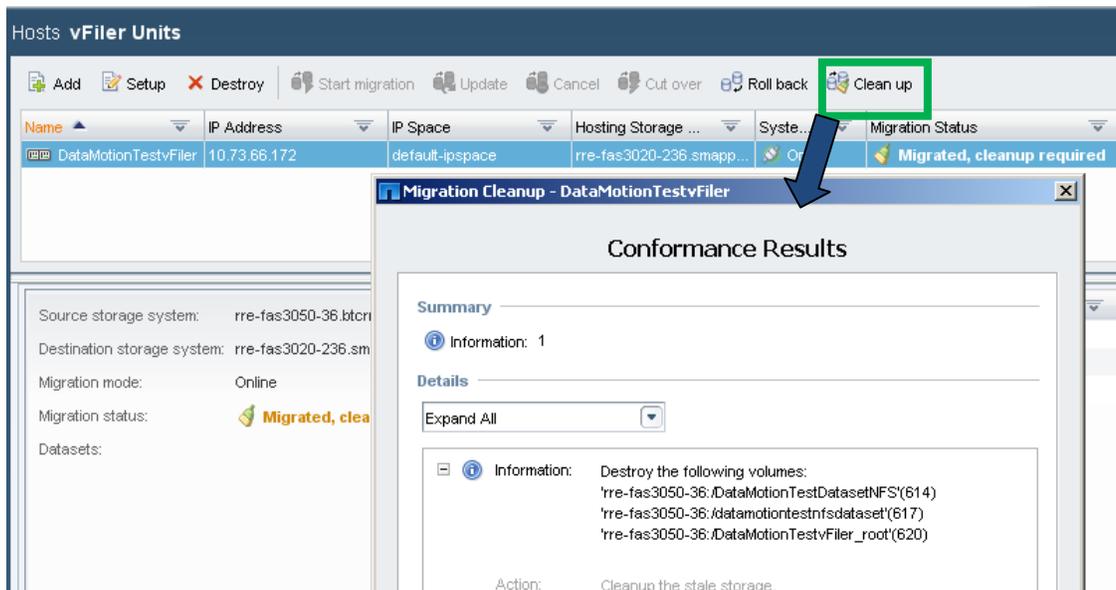


Figure 39) Migration cleanup (1).

2. After cleanup, the storage system settings changes as shown in the Figure 40:

The screenshot displays the NetApp vFile Units management interface. At the top, a toolbar includes actions like Add, Setup, Destroy, Start migration, Update, Cancel, Cut over, Roll back, and Clean up. Below this is a table of vFile Units. The table has columns for Name, IP Address, IP Space, Hosting Storage, System status, and Migration Status. A single entry is visible: DataMotionTestvFiler, with IP Address 10.73.66.172, IP Space default-ip-space, Hosting Storage rre-fas3020-236.smapp..., System status Online, and Migration Status Not started. The Migration Status cell is highlighted with a green box. Below the table, there are two sections: General and Hosting storage system settings. The General section shows Protocols (CIFS, NFS, iSCSI), Mirrored (Yes), Backup destination (No), and Backup source (No). The Service status section shows NFS (Up), CIFS (Down), and iSCSI (Up). The Hosting storage system settings section, also highlighted with a green box, shows Host name (rre-fas3020-236), IP address (10.73.66.236), System status (Online), Login credentials status (Good), NDMP status (Up), and NDMP credentials status (Good).

Name	IP Address	IP Space	Hosting Storage ...	Syste...	Migration Status
DataMotionTestvFiler	10.73.66.172	default-ip-space	rre-fas3020-236.smapp...	Online	Not started

**General**

Protocols: CIFS, NFS, iSCSI

Mirrored: Yes

Backup destination: No

Backup source: No

**Service status**

NFS: Up

CIFS: Down

iSCSI: Up

**Hosting storage system settings**

Host name: rre-fas3020-236

IP address: 10.73.66.236

System status: Online

Login credentials status: Good

NDMP status: Up

NDMP credentials status: Good

Figure 40) Migration cleanup (2).

## APPENDIX B: NETWORK CONFIGURATION BEST PRACTICES FOR DATA MOTION

Properly configuring the storage network is of utmost importance for successful and secure execution of Data Motion. This section describes an example setup that has been implemented in accordance with the Data Motion networking best practices.

Figure 41 illustrates a sample setup. There are two high-availability pairs: FAS1 and FAS2, and FAS3 and FAS4. The vFiler unit called vFiler1 running on FAS1 needs to be migrated to FAS3 by using Data Motion.

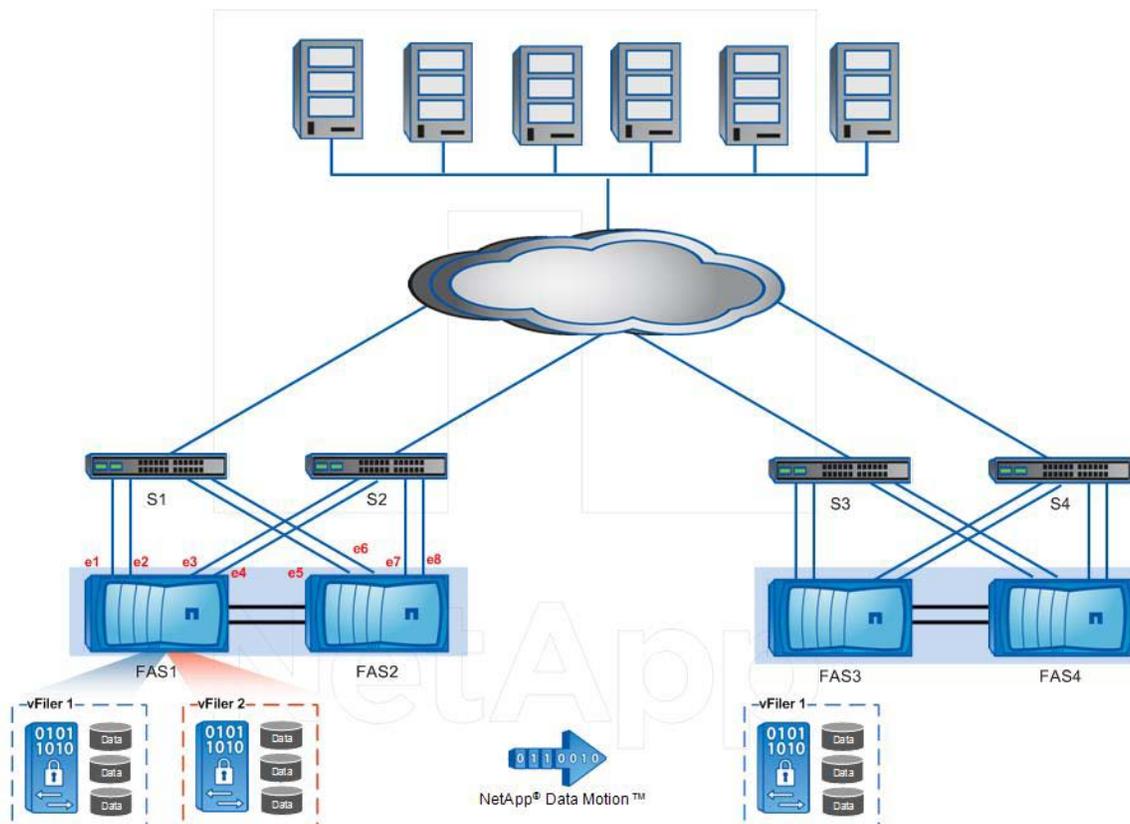


Figure 41) Network setup.

Table 14 describes the recommended method of configuring the networking on the storage systems and the network switches for NetApp Data Motion.

Table 14) Steps to configure networking on the storage system and network switches for Data Motion.

Step	Procedure
1	Configure VIF on the storage controller.
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 IP address to the appropriate VLAN interface over the VIF.
8	Add default route to the vFiler.

### STEP 1: CONFIGURE VIF ON THE STORAGE CONTROLLER

For simplicity, Figure 42 shows the connection between just one NetApp storage controller and the network switches. However the same configuration has to be repeated for all the storage controllers and the corresponding switches.

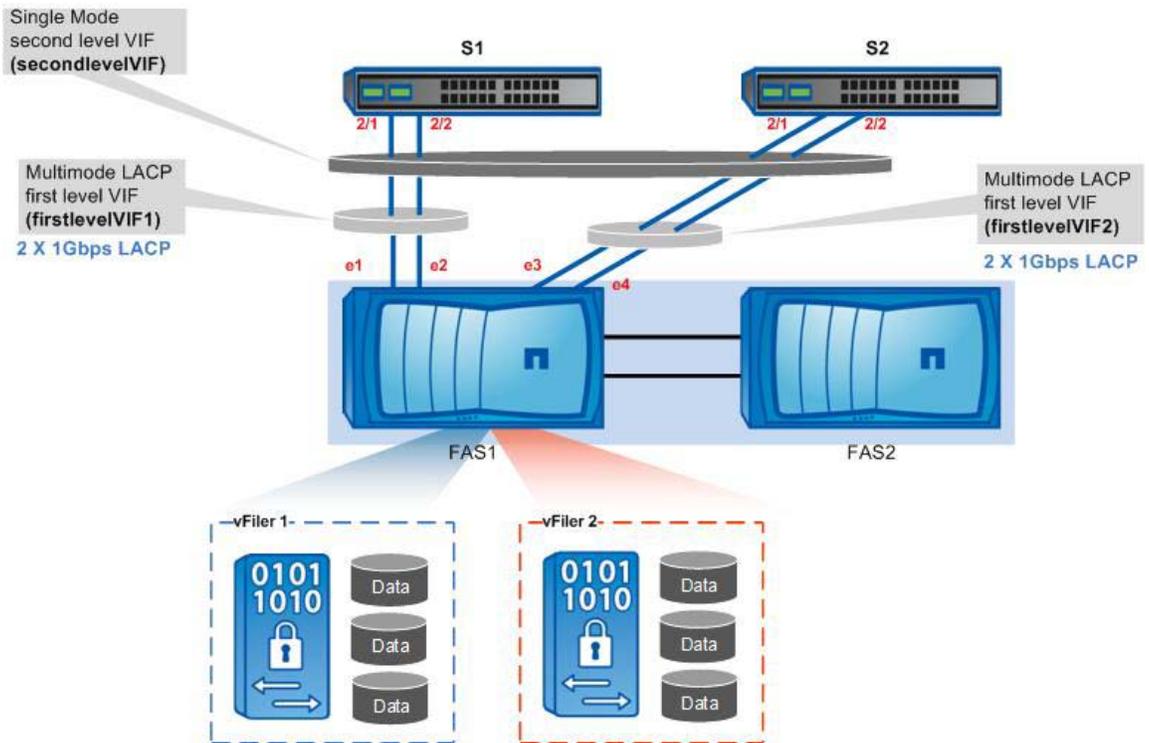


Figure 42) Connection between network switches and a NetApp storage controller.

In this example, the four 1Gbps links from the storage controller to two network switches have been configured as a multilayered VIF. e1 and e2 form a first-level multimode LACP VIF, 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: CREATE 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: ENABLE 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 ports to a VLAN trunked interface to support the VLAN defined.

In the switch S1, the configuration may 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 NetApp 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: CREATE IPSPACE

Create the IPspace where the vFiler unit will reside:

```
ipspace create ipspace1
```

### STEP 5: ASSIGN THE APPROPRIATE VLAN INTERFACE OVER THE VIF TO THE IPSPACE

To assign an interface to the IPspace, 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
```

Now assign the interface to the IPspace. In the following example, the single mode VIF is assigned to the IPspace created in step 1.

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

Check the available IPSpaces and the corresponding interfaces.

```
ipspace list
```

**STEP 6: 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: ASSIGN IP ADDRESS TO THE APPROPRIATE VLAN INTERFACE OVER THE 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: ADD DEFAULT ROUTE TO THE VFILER UNIT**

Add the default route to the vFiler unit:

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

Now the setup looks like Figure 43.

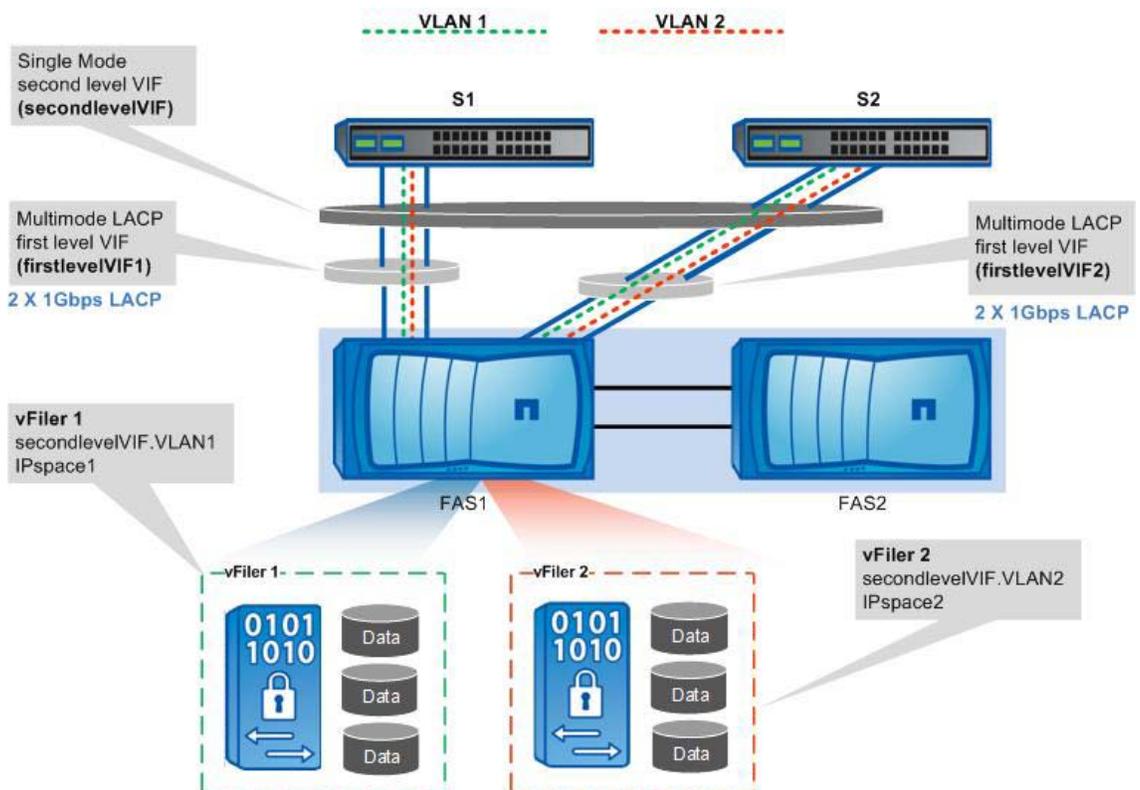


Figure 43) Setup.

**Note 1:** 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
```

Add the default route to the second vFiler unit:

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

**Note 2:** 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 1
```

**Note 3:** 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 10Gbps single mode VIF for Cisco Catalyst switch
- ✓ 2 X 10Gbps multimode VIF for Cisco Nexus switch running VPC

Steps to enable VLAN on the switch vary depending on the type of network switch selected. For a Cisco Nexus switch, following 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 that Cisco Nexus 5000 does **not** require an administrator to identify the VLAN trunk encapsulation type of 802.1q; it is enabled by default.

## REFERENCES

1. [Data ONTAP 7.3.3 MultiStore Management Guide](#)
2. Provisioning Manager 4.0 Workflow Guide
3. SnapMirror [Asynchronous](#) and [Synchronous](#) Best Practices
4. Non-disruptive Operations Community - <http://communities.netapp.com/groups/non-disruptive-operations>

## AUTHORS

[Larry Touchette](#), Technical Marketing Engineer, NetApp

[Roger Weeks](#), Technical Marketing Engineer, NetApp

[Preetom Goswami](#), Technical Marketing Engineer, NetApp

NetApp welcomes any comment or suggestions on this document. You can send your comments directly to the authors or post it in the '[Non-disruptive Operations Community](#)'.

## ACKNOWLEDGEMENT

Hamish McGovern, NetApp

Paul Feresten, NetApp

Bill L Heffelfinger, NetApp

Brian Gracely, NetApp

Uday Shet, NetApp

Jyoti Kurmi, NetApp

Nagender Somavarapu, NetApp

Devender Mudhiganti, NetApp

Srinivas Hemige, NetApp

Kaushik Sarkar, NetApp

Suresh Kumar, NetApp

Sridhara Gangoor, NetApp

Sitakanta Chaudhury, NetApp

Ratheesh Ramachandran, NetApp

Diyuan Liu, NetApp

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein must be used solely in connection with the NetApp products discussed in this document.

