

# Virtual Data Movers on EMC VNX

## Abstract

This white paper describes the high availability and portable capability of the Virtual Data Mover (VDM) technology delivered in the EMC® VNX™ series of storage platforms. VDMs are delivered standard with the VNX Operating Environment.

March 2016

Copyright © 2016 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on [EMC.com](http://EMC.com).

VMware is a registered trademark or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h10741.3

## Table of Contents

<b>Executive Summary .....</b>	<b>5</b>
<b>Audience.....</b>	<b>5</b>
<b>Technology Introduction .....</b>	<b>5</b>
Physical Data Mover versus VDM.....	6
<b>VDM States and Functionalities .....</b>	<b>8</b>
Creating a VDM.....	8
Changing the VDM State.....	9
Moving a VDM within the same system.....	9
<b>Planning Considerations.....</b>	<b>10</b>
VDM Names.....	10
VDM root file system size .....	10
Internationalization modes.....	10
Backing up VDMs with NDMP .....	11
Antivirus.....	11
Name resolution when moving VDMs .....	11
<b>VDM Replication.....</b>	<b>12</b>
VNX Replicator.....	12
VDM MetroSync .....	13
RecoverPoint .....	14
MirrorView.....	14
Symmetrix Remote Data Facility.....	14
Replicating a VDM for NFS with ESX datastore .....	14
<b>VDM Migration .....</b>	<b>15</b>
System Configuration Migration .....	16
Data Mover Level Configuration Migration Information .....	17
Cabinet Level Configuration Migration:.....	18
VDM Migration .....	19
Things that will be migrated during a VDM level migration:.....	20
Limitations:.....	21
File System Level Migration .....	21
Things that will be migrated during a file system level migration: .....	22
Limitations:.....	23
Troubleshooting .....	23
<b>Use cases and test results .....</b>	<b>24</b>
Disaster Recovery with Replication, Checkpoints, and VDMs .....	24
Boston .....	24
Miami.....	25

Interoperability.....	27
Conclusion.....	27
References.....	27

## Executive Summary

A Virtual Data Mover (VDM) is an EMC® VNX™ software feature that enables the grouping of Common Internet File Systems (CIFS) and/or Network File Systems (NFS) environments and servers into virtual containers. By using VDMs, it is possible to isolate CIFS and/or NFS environments from each other making them more secure, easier to replicate, and easier to migrate.

## Audience

This white paper is intended for storage and networking administrators, customers, and other users that are looking for an overview of the standard VDM functionality that is included in VNX series storage systems.

## Technology Introduction

VDMs contain the data needed to support one or more CIFS and/or NFS servers and their file systems. Each VDM has access only to the file systems mounted to that VDM, providing a logical isolation between physical Data Movers and other VDMs on the VNX system. When a VDM is created, a 128 MB in size root file system is created for that VDM. Only one root file system is created per VDM. This is the file system that stores the CIFS and/or NFS identity information. Data file systems are mounted to mount points created on the VDM root file system, and user data is kept in those data file systems.

Implement VDMs to:

- Enable replication of segregated CIFS/NFS environments.
- Simplify migration of segregated CIFS/NFS environments.
- Isolate CIFS servers and/or NFS servers to provide a higher level of security. This is particularly valuable in multitenant environments.
- Separate VLANs and file system mounts for different home directories, thus isolating home directory databases and their associated users.

Accessing data from a VDM is no different from accessing data that resides on a physical Data Mover.

The multi naming domain solution implements an NFS server per VDM named 'NFSendpoint'. The VDM acts as a container that includes the file systems exported by the NFS endpoint or the CIFS server, or both. These file systems on the VDM are visible through a subset of the Data Mover network interfaces attached to the VDM. The same network interface can be shared by both CIFS and NFS protocols on that VDM. Customers should expect no performance impact. VDMs perform in the same way as the physical Data Mover.

VNX provides a multi-naming domain solution for the Data Mover in the UNIX environment by enabling the implementation of an NFS endpoint(s) per VDM. The Data Mover hosting several VDMs is able to serve the UNIX clients that are members of different Lightweight Directory Access Protocol (LDAP) or Network Information Service (NIS) domains, assuming that each VDM works in a unique domain namespace. Similar to the CIFS service, several NFS domains are added to the physical Data Mover to provide access to the VDM for different naming domains. Each of these NFS endpoints is assigned to one or more Data Mover network interfaces. Any number of different Domain Name System (DNS) or Lightweight Directory Access Protocol (LDAP) domains may be specified.

The `server_nsdomains` command is used to select which domain you would like to specify on a per VDM basis.

### Physical Data Mover versus VDM

A Data Mover is a component that runs its own operating system (OS). This OS retrieves data from a storage device and makes it available to a network client. The Data Mover can use the NFS, CIFS or pNFS protocols.

The configuration and control data is stored in the root file system of a physical Data Mover, and includes:

- Databases (including localgroup.db and homedir, CIFS shares, and Registry directories).
- VDM root file systems (which store the VDM configuration files).
- Configuration files (passwd, group, viruschecker.conf, Group Policy Object (GPO) cache, and Kerberos file).
- Audit and event log information.
- Data file system mount points (appear as directories).
- BranchCache.

All VDM CIFS event logs are stored in the global event log of the physical Data Mover.

Each VDM stores its configuration information in its respective VDM root file system. The VDM configuration file is a subset of the configuration files on the physical Data Mover.

In addition, there are operations that can only be performed at the physical Data Mover level. Operations performed at the Data Mover level affect all VDMs.

These operations include:

- Stop, start, or delete services (CIFS, MPFS, viruschk)
- Data Mover failover
- Parameter changes

Figure 1 shows the VDM configuration within the physical Data Mover.



Figure 1: VDM root file systems within a physical Data Mover

Table 1 lists the location of certain configuration files relevant to a physical Data Mover and a VDM.

Table 1: Data Mover configuration files

Configuration Files	Physical Data Mover	Virtual Data Mover
Default CIFS database ✓		
Security ✓		
Networking ✓		
Internationalization mode ✓		
Virus checker ✓		
Parameters ✓	✓	
Standby Data Mover assignment and failover policy ✓	✓	
Local group ✓	✓	
HomeDir ✓	✓	
Share and repository directories		
File systems		
Kerberos		

A global (default) CIFS server and those CIFS servers within a VDM cannot coexist on the same Data Mover.

- A global, or default, CIFS server is assigned to all interfaces and grants access to all shares on that Data Mover.
- CIFS servers within a VDM require specific interfaces to be assigned to them which will grant access to those shares on that VDM.

If a VDM exists on a Data Mover, a CIFS server must be created on that VDM to grant access to all CIFS shares on that VDM.

The maximum number of VDMs per VNX array corresponds to the maximum number of file systems per Data Mover, which is 2,048. Realistically, one would not create a VDM without populating it. The objects that reduce the file system count on a system are as follows:

- The root file system of a VDM.
- Any file system created on the DM or VDM.
- A checkpoint of a file system.
- A storage pool as shown in the `nas_pool -list` command.

## VDM States and Functionalities

Table 2 references the different states of a VDM.

**Table 2: VDM states and functionalities**

VDM state	CIFS active	User file systems	VDM reloaded at boot	Failover	Failback
Loaded	Yes	Accessible	Yes	No	Yes
Mounted	No	Inaccessible	No	Yes	No
Temporarily unloaded	No	Inaccessible	Yes	No	No
Permanently unloaded	No	Inaccessible	No	No	No

VDM states can be changed with the `nas_server` command. These states are discussed in the following sections.

### Creating a VDM

The default state upon VDM creation is loaded. In the loaded state, the VDM is fully functional and active. A VDM must be in the loaded state to allow most configuration changes, such as the addition of CIFS servers or NFS endpoints.

A VDM can only be loaded on one physical Data Mover at a time. Meaning, you cannot have a VDM in a loaded state on two physical Data Movers within a VNX system at the same time. If you need to move the VDM, the network interfaces used by its CIFS servers must be available on the destination Data Mover. VDMs within a VDM are not supported.



Before VDM creation, consider the following:

- If you do not specify the state, the VDM is created in the loaded (default) state.
- If you do not specify a name for the VDM or its root file system, a default name is assigned. See [Planning](#) for more information about VDM naming conventions.
- The VDM root file system is created from an existing storage pool.
- The default size of the root file system is 128 MB.
- During the planning and design phase, implement VDMs before you construct the environment.
- Network configuration should be considered. For example, NTP, DNS, and domain controllers.
- While designing VDMs, create VDMs with file access in mind. Administrators should consider load balancing the VDMs in a given physical Data Mover.

Administrators should consider the following before creating, configuring, or managing VDMs:

- Establish a naming convention that indicates the function and easily identifies the VDM and its root file system.
- Allocate appropriate space for a VDM and its configuration file system. The default size is 128 MB.
- Consider FS groupings and which CIFS servers and/or NFS endpoints reside on which physical Data Mover.

## Changing the VDM State

Unlike a state change from loaded to mounted, a state change from loaded to unloaded (tempunloaded or permunloaded) requires the VDM to have no mounted file systems. The VDM root file system is not deleted and is available to be mounted. This transition to the unloaded state is required when you want to stop all activity on the VDM, but do not want to delete the VDM root file system with its configuration data.

Changing the VDM state from loaded to mounted, tempunloaded, or permunloaded shuts down the CIFS servers and/or NFS endpoints on the VDM, and makes the file systems inaccessible to the clients through the VDM. When you unload a VDM from a physical Data Mover, you must specify the intent for the unload operation: permanent or temporary.

## Moving a VDM within the same system

When you move a VDM from one physical Data Mover to another, the Control Station performs the following operations:

- Unloads the VDM from the source Data Mover.

- Unmounts all file systems on that VDM.
- Loads the VDM and mounts all file systems on the target Data Mover.

Before you move a VDM, consider the following:

- For every CIFS server in a VDM, ensure that the CIFS server's interfaces are available and identically named on the physical Data Mover to which you move the VDM.
- The amount of time it takes to update the network information needed to make the VDM available after you move the VDM to a physical Data Mover with a different IP address. Network information might include, but is not limited to, static files, DNS, or the Windows Internet Naming Service (WINS).

## Planning Considerations

### VDM Names

When naming a VDM, it is recommended that you name it according to its function. This will help identify the VDM and its root file system. For example, you might group Marketing CIFS servers together into a VDM named Marketing so that the VDM is easily identifiable. The VNX assigns the VDM root file system a name in the form of `root_fs_vdm_<vdm name>`. For example, the root file system for a VDM named Marketing will be `root_fs_vdm_Marketing`. (If you name a VDM `vdm_Marketing`, the VNX does not duplicate the VDM part of the name. The root file system is still named `root_fs_vdm_Marketing`.)

If you rename a VDM, its root file system is renamed accordingly. For example, if you rename a VDM to HR, its root file system name changes to `root_fs_vdm_HR`.

### VDM root file system size

A root file system is assigned to a VDM when it is created. The default size is 128 MB, which is the same for a physical Data Mover. In an environment with a large number of users or shares, you might need to increase the size of the root file system. You cannot have the system extend the root file system automatically.

### Internationalization modes

The VNX supports clients in environments that use multibyte character sets. When supported, multibyte character sets enable universal character encoding standards (Unicode).

When a VDM is created, its Internationalization mode is set to the same mode as the Data Mover in which it resides. When the VDM is unloaded, its mode matches the last physical Data Mover on which it was loaded.

## Backing up VDMs with NDMP

A full path is required to back up VDM-mounted file systems with NDMP backup and the `server_archive` command. An NDMP example is

`/root_vdm_Marketing/fs`. A server archive example is:

```
server_archive <movername> -w -f /dev/clt410/ -J  
/root_vdm_Marketing/fs.
```

## Antivirus

In addition to CIFS servers created within a VDM, a global CIFS server is required for antivirus functionality. A global CIFS server is created at the physical Data Mover level. There must be at least one CIFS server on the physical Data Mover if you use an Antivirus solution.

## Name resolution when moving VDMs

Consider name resolution when moving a VDM to another physical Data Mover and using different IP addresses. The name resolution method depends on whether the environment uses DNS, WINS, or a static file.

### DNS

Consider DNS configurations on the server and client side when moving VDMs:

- The destination Data Mover DNS resolver must be correctly configured to satisfy the DNS resolution requirements for the VDM. (see `server_dns` command).
- The VDM load operation updates the DNS servers configured on the physical Data Mover.
- The administrator can force DNS database updates to all DNS servers to which the clients might be pointing.
- When you move the VDM to the destination Data Mover, the DNS server automatically updates and displays the correct entry.
- The update time between DNS servers depends on how the DNS servers are configured.
- All the name resolutions for the VDM are confined in the VDM if the VDM name server domain configuration is enabled (see `server_nsdomains` command).

### WINS

When moving VDMs in an environment using WINS, consider the following:

- The VDM load operation updates WINS servers configured on the VDM.
- The update time between WINS servers depends on how WINS servers are configured.

- Administrators can force WINS database updates to all WINS servers to which the clients might be pointing.
- To clear and update the WINS cache for clients, use one of the following methods:
  - Restart the computer.
  - Wait until the cache is cleared (TTL 10 minutes).
  - From the DOS prompt, run the command `nbtstat -R` (preferred method).

### Static file (LMHOSTS/hosts)

The static file (hosts) used by clients must be updated with the IP address of each new CIFS server.

### Limitations

The following limitations currently apply to VDMs:

- VDMs support CIFS and NFS protocols over TCP. All other file protocols such as iSCSI, FTP, SFTP, and FTPS are not supported. The NFS clients must support NFSv3 or NFSv4 over TCP to connect to a NFS endpoint.
- The VDM for NFS multidomain feature needs to be administered by using the CLI, the GUI is not supported.
- IP replication failover support for Local Groups must include VDMs.
- All VDMs share the Replay cache (XID cache).

## VDM Replication

VDMs can be individually replicated by using either VNX Replicator or VDM MetroSync. Also, RecoverPoint, MirrorView, and SRDF technologies can be leveraged for cabinet-level DR solutions. Depending on the replication technology you are using, the task of replicating these file systems may differ. Please review the EMC VNX Replication Technologies whitepaper for additional information.

When performing replication, it is important to understand that in order for a CIFS or NFS environment to be fully functional and accessible on a remote VNX system, the complete CIFS or NFS working environment must be replicated. This will include the VDM and its configuration along with any configuration built exclusively on the physical Data Mover such as DNS, NIS, NTP, Local passwd and group, Usermapper client, FTP/SFTP, LDAP, HTTP, CEPP, CAVA, Server Parameters, netgroup, nsswitch, Hosts, and so on. In addition to the environment configuration, the file systems associated with the CIFS or NFS environment must be replicated.

### VNX Replicator

VNX Replicator is a file level, IP based replication solution. VNX Replicator is capable of replicating at the VDM and file system granular level. VNX Replicator is an asynchronous replication solution. For additional information on VNX Replicator,

please consult the *Using VNX Replicator* document. Additional information for VNX Replicator can be found on EMC Online Support.

## VDM MetroSync

VDM MetroSync is a Disaster Recovery (DR) solution for VNX2 File which leverages a MirrorView/S replication session to create a zero data loss replication solution at a VDM granularity. It allows for replication of a VDM along with all of its contents including file systems, checkpoints, checkpoint schedules, CIFS servers, exports, interfaces, and so on. It can be configured in either an active/passive configuration where the active VDMs are constrained to one site, or an active/active configuration where each site has its own set of active VDMs. VDMs can be moved or failed over from one system to another as needed. The VDM MetroSync solution is shown in Figure 2.

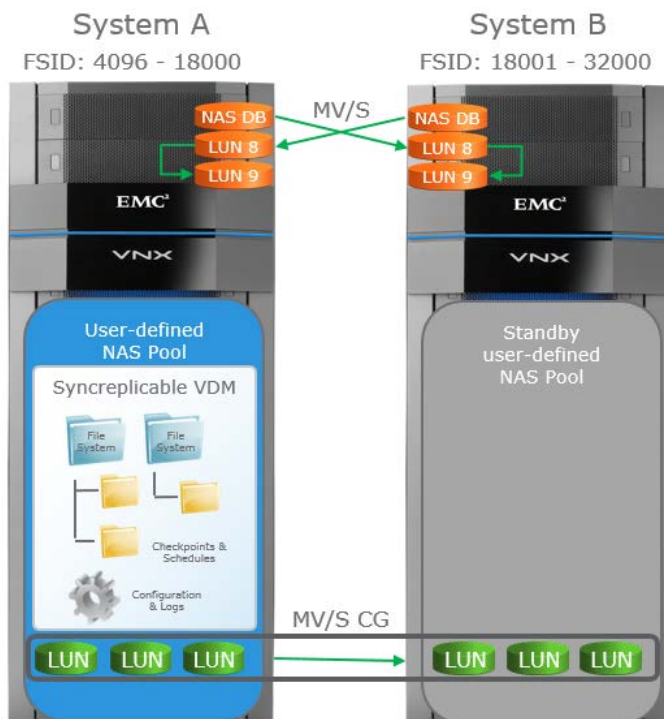


Figure 2: VDM MetroSync

VDM MetroSync Manager is optional software that can be installed on a Windows server which works with VDM MetroSync. It provides a GUI interface to display VDM MetroSync session information and run operations to move, failover, or restore VDMs. It also has the ability to continuously monitor sessions and automatically initiate failover when issues are detected.

With synchronous replication enabled between two systems, it is also possible to add asynchronous replication to a third system by using Replicator. This allows the third system to be located further away and enables it to be used as a backup and recovery solution. When VDMs are moved or failed over between the VDM MetroSync systems, the Replicator sessions to the third system are preserved. Since the Replicator

checkpoints are replicated along with the VDM, a common base checkpoint is available which removes the requirement for a full synchronization after failover. The Replicator sessions can be incrementally updated and restarted on the new system where the VDM is active.

For more information on VDM MetroSync, please refer to the *VDM MetroSync for VNX2* white paper, available on the EMC Online Support website at <https://support.emc.com>.

## RecoverPoint

RecoverPoint is a block level, iSCSI or Fibre Channel based replication solution. RecoverPoint can be leveraged to replicate NAS solutions at the cabinet-level. For additional information regarding RecoverPoint technology, please consult the RecoverPoint product page on EMC Online Support.

## MirrorView

MirrorView is a block level, iSCSI, Fibre Channel, or FCoE based replication solution. MirrorView can be used to replicate a VDM environment at the cabinet-level. For additional information, please consult the MirrorView Knowledgebook for Releases 30 – 33 whitepaper. Additional information can be found on EMC Online Support.

## Symmetrix Remote Data Facility

Symmetrix Remote Data Facility (SRDF) is a block level, IP, fibre channel, or ESCON (Enterprise Systems Connection) based replication solution. SRDF can be used to replicate a VDM environment at the cabinet-level. For additional information, please consult the SRDF product page on EMC Online Support.

## Replicating a VDM for NFS with ESX datastore

When replicating VDM for NFS file systems used for ESX datastores, use the following procedure to ensure that the virtual machines (VMs) do not go offline:

---

**Note:** Follow these steps exactly in the given order or the ESX server will lose datastore access. If the VDM is failed over before the user file system, the ESX server receives an error NFS3ERR\_STALE (Invalid File Handle). The ESX client then considers the file system to be down.

---

This procedure assumes that your NFS datastore is connected to VNX through an IP address and not with a name that has to resolve in DNS.

1. Fail over the Production File System mounted on the VDM.
2. Down the Interface on the source.
3. Fail over the VDM.
4. Up the Interface on the destination.

To avoid an error, when the Production File System mounted on a VDM fails over, the VDM interface is set down using the `server_ifconfig` command on the source. The VDM is then failed over. As the VDM restarts on the replication site, the VDM interface can be set up using the `server_ifconfig` command.

## VDM Migration

VDM migrations leverage VNX Replicator technology and can be performed using the VNX File OE Command Line Interface (CLI) or the VNX File Migration GUI Tool. Figure 3 shows a screenshot of the VNX File Migration GUI Tool.

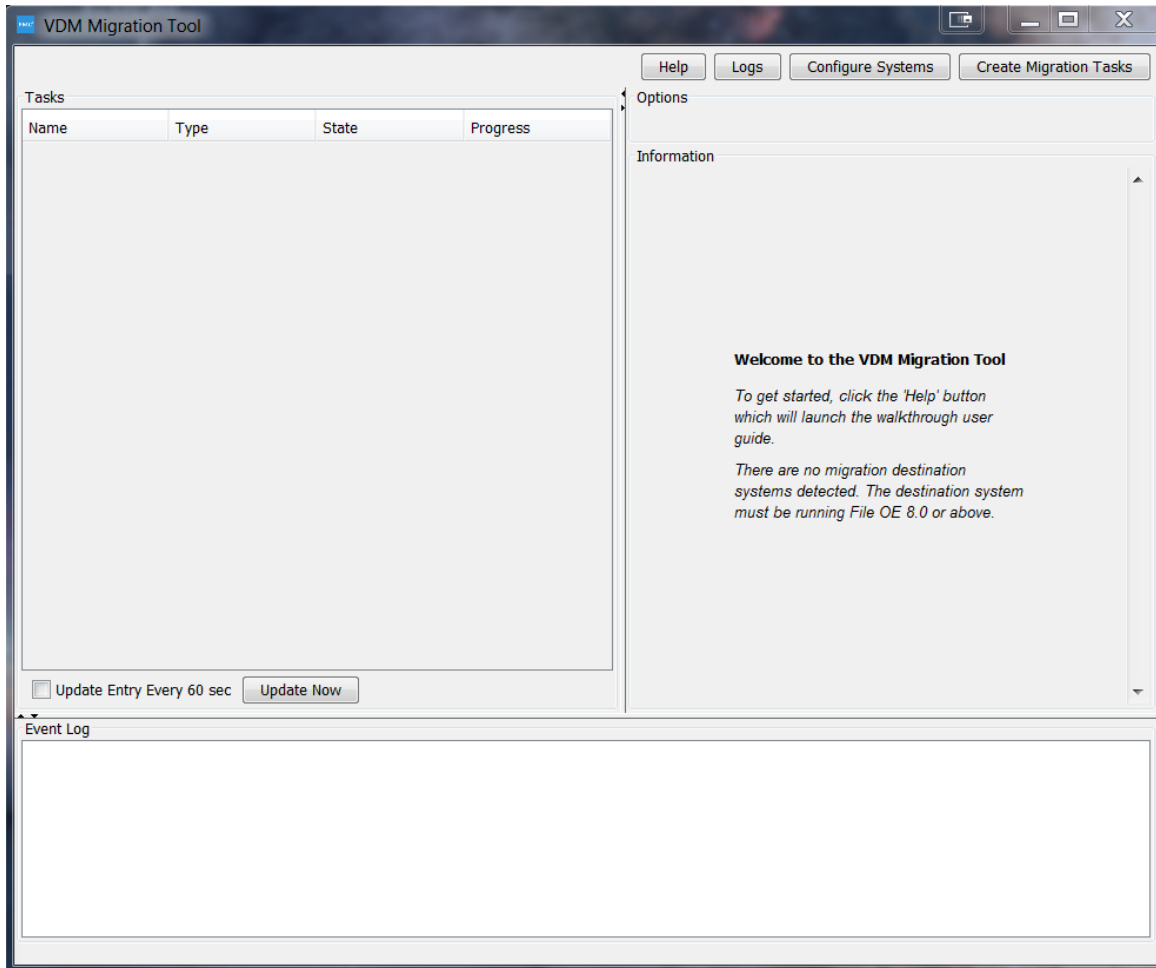


Figure 3: screenshot of the VNX File Migration GUI Tool

VNX systems on code version 8.1.6.96 / 05.33.006.5.096 and later have the ability to perform VDM migrations and Data Mover (DM) configuration migrations with the option to preserve the Disaster Recovery solution (the VNX Replicator replication sessions). This is beneficial for customers because it simplifies migrations and allows customers to maintain their Disaster Recovery solution throughout the migration process. This prevents the need to rebuild the Disaster Recovery solution manually and perform a full copy of the data. The migration commands must be run from the



target VNX system and will not be back ported to pre-VNX File OE 8.1 systems. It is not guaranteed that the file systems will retain their file system IDs when migrating to the destination system. If the file system ID changes, the NFS clients will need to remount the exports.

In order to leverage VDM migration, the source and the destination VDM/FS must be located on different storage systems. Migration between DMs in the same storage array is not supported at this time. Moving VDMs across DMs can be accomplished by using the `nas_server -vdm <vdm_name> -move` command or the Unisphere GUI.

The maximum number of file systems that can be migrated concurrently is 256 per DM. If the number of file systems exceeds the restrictions, the system will put them in a queue and migrate them in sequence. Remember that the root VDM file system is counted as a file system.

For additional information regarding these workflows, please consult the 'Using VNX File Migration Technical Notes' document.

You can perform these operations using the CLI or the VNX File Migration Tool which is a standalone GUI. The GUI can be downloaded from EMC Online Support by searching VNX File Migration Tool. Complete migrations can be done using CLI or the GUI tool.

For additional information on how to use the GUI click the Help button and it will launch a step by step guide.

## System Configuration Migration

The system configuration migration command is implemented to facilitate VDM Level Migration. This command must be initiated from the destination side of the migration. The system configuration can be migrated using the `migrate_system_conf` command.

The following configurations on a specific physical Data Mover can be migrated using this command. The administrator can pick and choose which to migrate if all are not needed. The VDM configurations are not included in the list as they will be migrated along with VDM.

- Data Mover Services
  - DNS
  - NIS
  - NTP
  - passwd and group files
  - Usermapper client (IP address of external Usermapper service)
  - FTP/SFTP/FTPS
  - LDAP
    - nsswitch.conf



- HTTP
- CEPP
- CAVA
- Server Parameters
- Netgroup
- Hosts
- Cabinet level services
  - Usermapper service

## Data Mover Level Configuration Migration Information

### DNS

By default, the DNS configuration on the source DM does not overwrite the destination's current DNS configuration, as you can have multiple DNS servers. You can however, choose the `-overwrite_destination` option to enforce a migration of all DNS (and other service) settings from Source to Destination. The customers can choose to configure DNS on the destination DM manually. In this case the configuration tool will not migrate the DNS configuration.

### NTP

The NTP client and time zone settings will be migrated from source Data Mover to destination Data Mover.

### Password and group files

The passwd and group files of source DM will replace the current passwd and group files of destination DM.

### LDAP

LDAP client settings will be migrated from source to destination.

- If `nsswitch.conf` exists on the source and LDAP is selected, `nsswitch.conf` will be copied to the destination.
- The LDAP client cannot be migrated if it is configured to use Kerberos to login in to the source side.

### NSswitch

The `nsswitch` file will be migrated from the source Data Mover to the destination Data Mover.

### CEPP & CAVA

The CEPP / CAVA configuration file will be migrated from the source Data Mover to the destination Data Mover.

- The local group on the source Data Mover and EMC virus-checking rights on that local group are not migrated to the destination.
- You need to reconfigure them through MMC on the destination manually after migration.

### Netgroup

The netgroup file will be migrated from the source Data Mover to the destination Data Mover.

### Hosts

The hosts file will be moved from the source Data Mover to the destination Data Mover.

### Server\_Parameters

The server\_parameters in the source Data Mover will be migrated from the source Data Mover to the destination Data Mover.

- Parameters added to the `/nas/server/slot_x/param` and `/nas/site/slot_param` file will be set on the destination machine.
  - This option will require a DM reboot after the parameters are migrated over to the destination side.

## Cabinet Level Configuration Migration:

### Usermapper Service

- If the destination system will be a replacement for the source system, you can choose to use the DM Configuration Migration command to migrate the Usermapper database which will work as the follows. There are a few scenarios to consider when migrating the Usermapper service: If the Usermapper service is primary on the source:
  1. Backup destination Usermapper DB if there is already an existing one.
  2. Copy the Usermapper DB from source to destination.
  3. Set the destination system to be the primary Usermapper.
  4. Change the source system to secondary mode pointing to the destination system.
- If the Usermapper service is secondary on the source:
  1. The destination will be configured as a secondary and point to the same primary the source system points to.
- If Usermapper is disabled on the source system:
  1. The system will prompt the user to confirm whether it is disabled on purpose.

2. The system will ask the user to disable the destination or enabled source manually and then retry the migration.

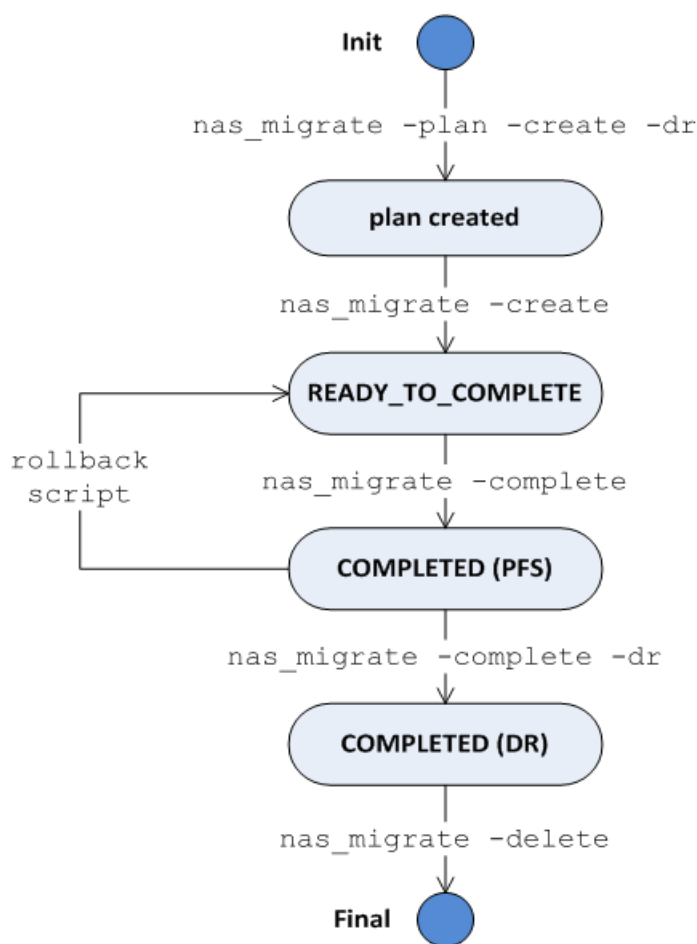
If the destination DM is configured with the Usermapper service, the administrator can overwrite the configuration if necessary.

- This overwrite case should only be used to redo a previously failed Usermapper migration.

## VDM Migration

VDM migration will migrate all the resources hosted by a specific VDM. This includes the VDM configuration and any file systems mounted to the VDM. VDM migrations are performed using the `nas_migrate` command issued from the destination of the migration. If possible, the file system IDs will be preserved on the destination system. IPv4 interfaces can be migrated if you would like the interface to remain exactly the same from source to destination, by using the `-take_over_ips` option. When not using the `-take_over_ips` option, the interfaces attached to the VDM on the destination must be created with the exact same name as the interface that is attached to the VDM on the source in order to execute a successful cutover of VDM migration. When manually creating interfaces be sure to note any other settings that might be required such as VLANs or MTU settings.

VDM migration only migrates the file systems and checkpoints detected during the migration plan creation. The system does not track file system mount/unmount or checkpoint create/delete/refresh operations triggered by an administrator or data services after the migration plan is created.



**Figure 4: VDM migration workflow**

The migration plan and create are online, without interruption to the users of the system. The complete (cutover) could result in a network outage. The cutover time depends on the number of file systems, the bandwidth of DM interconnect, and the throughput to the source file systems. See the *Using VNX File Migration Tech Note* for more information located on EMC Online Support.

**Things that will be migrated during a VDM level migration:**

- VDM root file system.
- All file systems mounted to the VDM.
- Read-only checkpoints mounted to this VDM.
  - Checkpoints associated with schedules that are in a paused or complete state will be migrated.
  - Checkpoints associated with schedules that are in an active or pending state will not be migrated.
- Checkpoint schedules will be migrated to destination side and started during the complete step.

## Limitations:

- Can only migrate one VDM per plan.
  - Cannot combine multiple VDMs in the same plan.
  - Multiple creates can be run simultaneously however bandwidth needs to be monitored carefully.
  - Cannot combine multiple plans in a create step.
- Cannot modify the plan after the create command has been initiated.
- Any checkpoints created on the VDM after the create command is initiated will not be migrated.
- Any Checkpoint(s) refreshed on the source will not be refreshed to the destination.
- Any Checkpoint(s) deleted from the source during the migration will not be deleted on the destination.
- Checkpoints refreshed on the source will not be refreshed to the destination.
- NDMP and Replication checkpoints are NOT migrated.
- Writeable checkpoints are NOT migrated.

## File System Level Migration

File system level migration will migrate one file system and the read only checkpoints mounted to it. The file system ID will be preserved if possible. This is also done using the `nas_migrate` command issued from the destination of the migration. If you would like to migrate a CIFS server that is used for a file system on the physical DM, you must move the file system and CIFS server to a VDM on the source system, and then migrate the VDM.

The destination file system will have the identical file system type, size, and FS log type as the source file system. The SavVol can be specified on the destination DM, which will allow the customer to specify a smaller SavVol in order to reduce space usage.

If creating the destination file system manually, then this destination FS must be empty, mounted to DM or VDM as read-only, have an identical size as the source FS and keep with any other destination restrictions of RepV2. The restrictions can be found in the *Using VNX Replicator* document on EMC Online Support.

The destination FS retains the same name as the source if possible. If there is a name conflict on the destination system, the system generates a new name such as “<source fs name>\_mig<source FSID>\_N” and verifies it is available on the destination system. A similar mechanism to file system name conflicts is used for any mount points created on the destination DM.

If the destination is a VDM, the system will try to create and use the same mount point name on the destination side. If the same mount point name is already used by another file system, the migration will fail.

The system cannot guarantee that the FSID will be preserved on the destination. The FSID of the file system will **NOT** be retained if there are FSID conflicts on destination system. The FSID of checkpoints will **NOT** be retained at all.

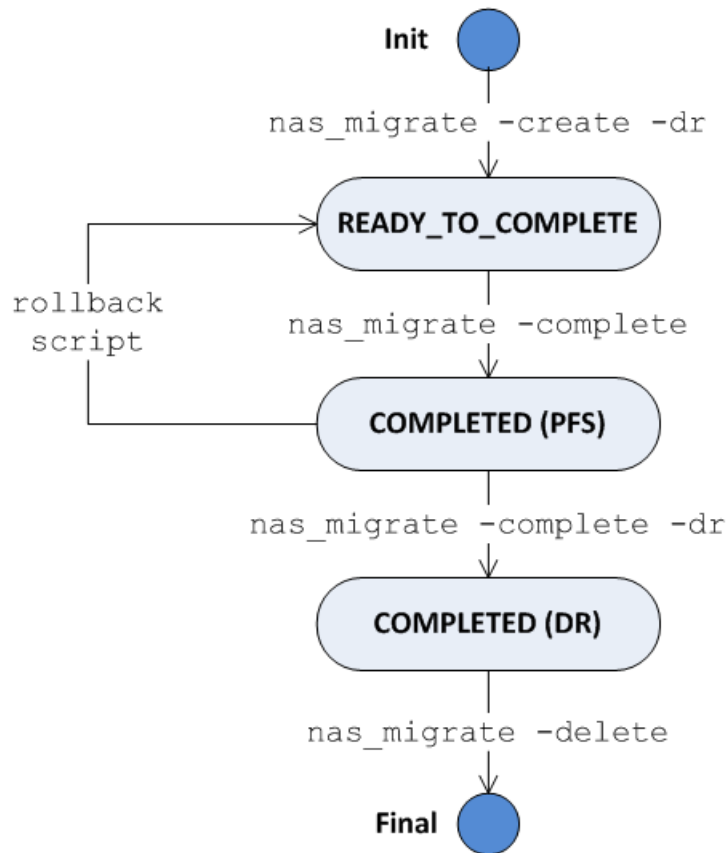


Figure 5: File System migration workflow

#### Things that will be migrated during a file system level migration:

- Specific file system.
- Read-only checkpoints for the file system.
  - Checkpoints associated with schedules that are in a paused or complete state will be migrated.
  - Checkpoints associated with schedules that are in an active or pending state will not be migrated.
- Checkpoint schedules will be migrated to destination side and started during the complete step.

## Limitations:

- Cannot combine multiple file systems in the same create command.
  - Multiple creates can be run simultaneously however bandwidth needs to be monitored carefully.
- Nested Mounted File Systems (NMFS) are NOT supported in FS Level Migration.
  - The component file systems of NMFS are supported.
- A Migration File System (MGFS) is NOT supported. Neither the source file system nor the destination file system can be MGFS. MGFS is a type of file system that might be seen in Celerra systems. You can determine file system type by looking at the properties of the file system.
- Any checkpoints created for this file system after the create command is initiated will not be migrated.
- Checkpoints deleted from the source during migration will not be deleted on destination.
- Checkpoints refreshed on the source will not be refreshed to the destination.
- NDMP and Replication checkpoints are not migrated.
- Writeable checkpoints are NOT migrated.

## Troubleshooting

While working with VDMs, an administrator may have to troubleshoot the following issues:

- **CIFS and or NFS server accessibility errors:** When no accessible CIFS or NFS servers are found, verify the VDM and VNX system logs (just as you would troubleshoot a physical Data Mover). Optionally, verify the state of the VDM. If the VDM is in the mounted state, change the VDM to the loaded state and verify again. In addition, check to ensure that the CIFS server is still joined to that domain and that the IP interface is up.

Perform these tasks to troubleshoot this error:

- Verify the CIFS protocol configuration.
  - Verify the attributes of a VDM.
  - Change the VDM state.
  - Verify the attributes of a VDM.
- **VDM unloading errors:** Do not unload a VDM with mounted file systems. Unmount file systems prior to unloading a VDM.

Perform these tasks to troubleshoot this error:

- Verify a list of mounted or temporarily unmounted file systems.
- Verify the attributes of a VDM.
- Unload a VDM temporarily.

## Use cases and test results

### Disaster Recovery with Replication, Checkpoints, and VDMs

The most common use case scenario for VDMs is described below:

Company ABC experienced a substantial growth in the demands for their services in a year and also projected a growth of approximately 40 percent over the next few years. To support this growth, the company planned to expand the business to a location outside of Boston (ABC's primary site). ABC established another data center in Miami (secondary site) and designated it to be the disaster recovery site.

Management realizes the importance of having a data recovery plan for their CIFS environment. They need to ensure that data is always available by asynchronously replicating over IP to the data center in Miami.

Company ABC has the following requirements:

- Leverage the Replicator, SnapSure™, and the VDM technology for the data recovery solution.
- Monitor the status and proper function of this replication to the destination site.
- Give the destination site read/write permissions to temporarily service the orphaned source-site users in the event of a disaster.
- Use the same source site after a failover when it becomes available.

Testing was done to ensure that data is always available by asynchronously replicating over IP to the data center in Miami.

### Boston

To support the disaster recovery plan, the administrator in Boston must first confirm that VDMs are in the loaded state. Most often, VDMs are in the loaded state because the CIFS shares or NFS exports on these VDMs are already active and data is accessible by all employees. Each server in a CIFS or NFS environment has associated file systems and supporting attributes.

In addition, the active checkpoint schedules on the source site created checkpoints to allow users and system administrators to perform file-level restores.

Figure 6 shows the configuration of the Boston primary site. The \\eng\_ne server is contained within the Engineering VDM and can see only the file systems for that specific VDM. There is only one root file system (root\_fs\_vdm\_Engineering) and the names for the source-user file systems are:



- Eng\_User
- Eng\_Shared

The share names for the source-user file systems are:

- User\_Files
- Shared\_Files

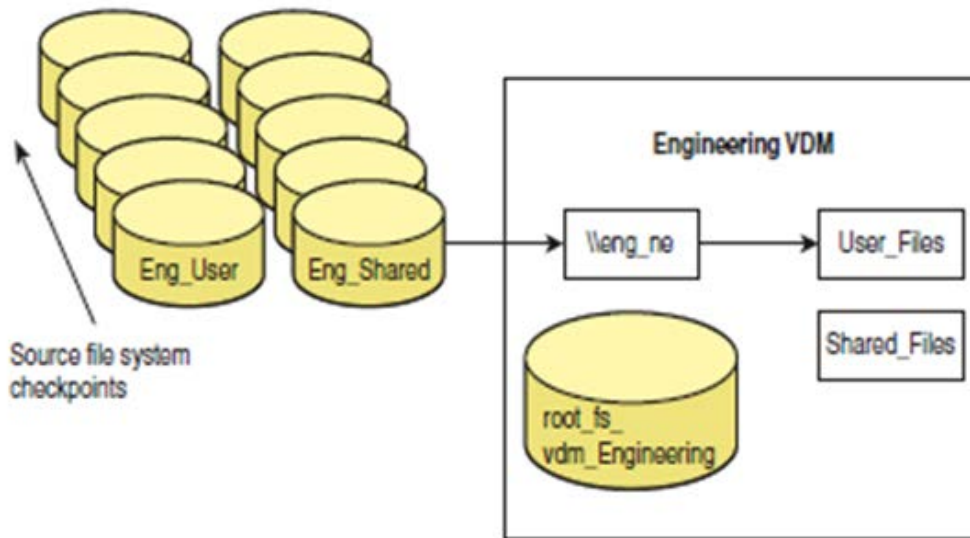


Figure 6: Configuration of primary site

### Miami

The VNX administrator needs to configure replication sessions for the Engineering VDM and the two file systems. Repv2 creates correctly sized and appropriately mounted VDMs and data file systems on the Miami VNX. Additionally, the VNX administrator needs to create and start the schedule for checkpoints that run at the same intervals as in Boston.

This architecture meets the company's requirements, which are as follows:

- The administrator has set up this disaster recovery replication to the destination site by replicating the VDM. This gives all users access to the CIFS servers and user file systems.
- Monitor the status and proper function of the replication pairs by monitoring the status of the VDM replication.
- In the event of a catastrophic disaster, the VNX administrator executes a Repv2 failover on all replication sessions from the Miami VNX, as seen in Figure 7 Repv2 then switches VDMs to a loaded state and remount data file systems in a read-write mode.

---

**Note:** The checkpoint schedule continues to run without the knowledge of the Repv2 failover.

---

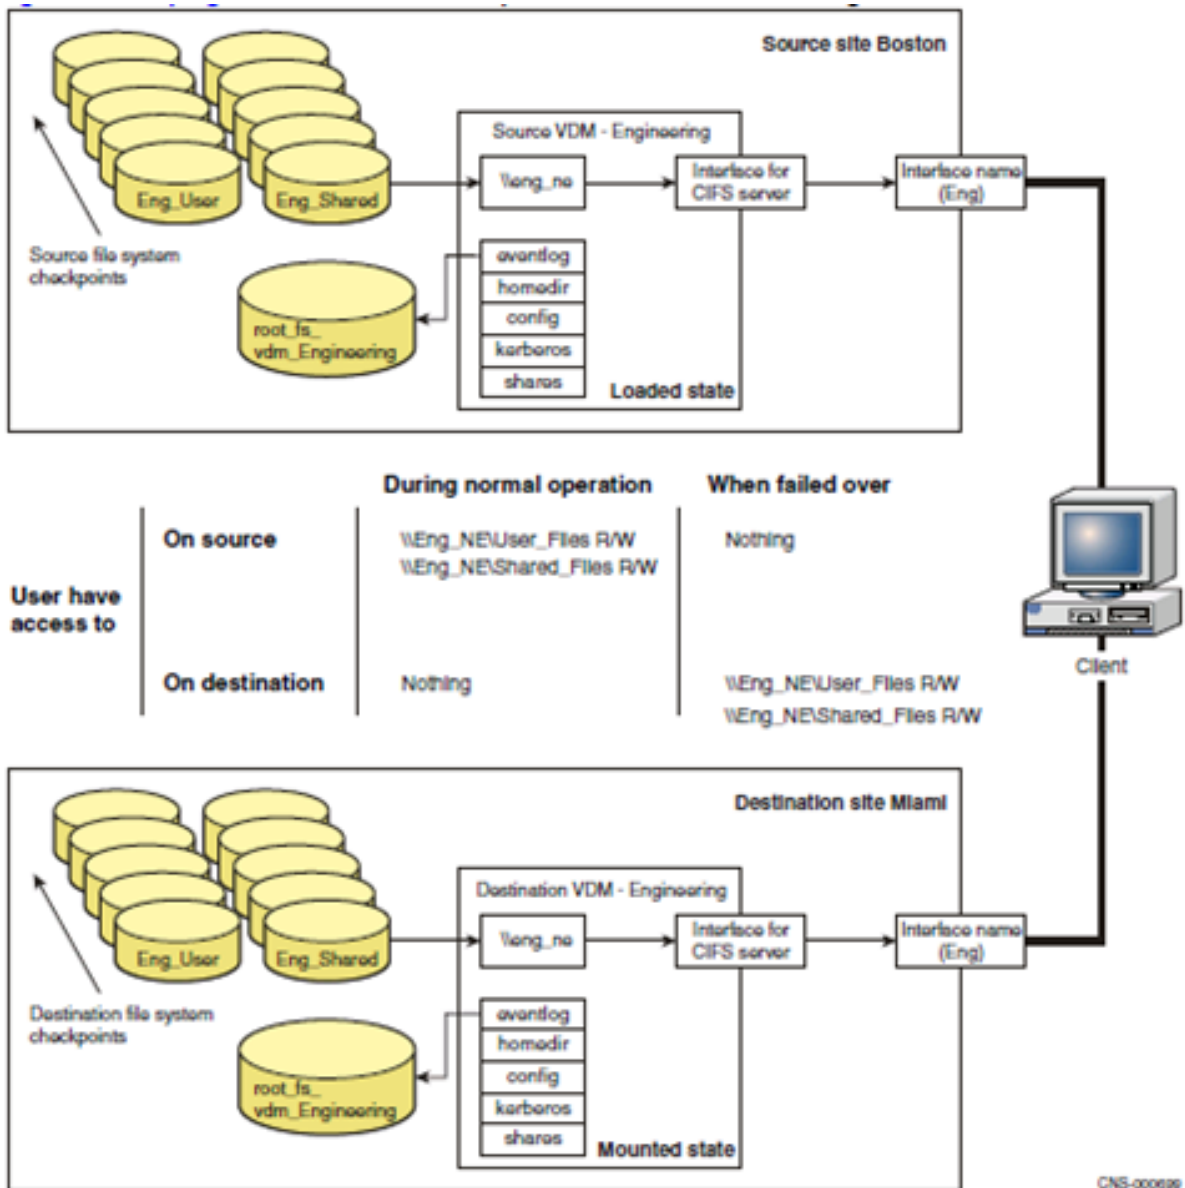


Figure 7: Disaster recovery replication

## Interoperability

VDMs interoperate with other VNX features as follows:

- After a VDM is created, CIFS servers and shares can be loaded into the VDM. Separating CIFS servers allows more flexibility to replicate content between CIFS servers.
- The VNX Home Directory feature can leverage VDM capabilities to replicate home directory environments between VNX cabinets. If the source site is lost in a disaster scenario or taken down for maintenance, the destination site can be brought up and all of the home directory services function as they did before the event.

## Conclusion

VNX enables users to effortlessly create and manage VDMs. While they require initial configuration, VDMs do not need daily management. They can greatly enhance storage system functionality. VDMs help simplify workload management by consolidating file systems and CIFS / NFS environments and servers into groupings so they can be replicated and migrated quickly and easily.

## References

For additional information the following documents can be located on EMC Online Support:

- *Configuring Virtual Data Movers on VNX*
- *EMC VNX Replication Technologies*
- *Using VNX Replicator*
- *Configuring NFS on VNX*
- *Configuring CIFS on VNX*
- *Configuring VNX User mapping*

To locate the VNX File migration Tool search EMC Online Support for VNX File Migration Tool. There you will find a zip that includes the standalone GUI tool and the Technical notes about the feature. There is also the help file in the tool that is a how-to on using the tool itself.