White Paper

STRATEGIES FOR MIGRATING BLOCK DATA TO EMC VNX SERIES UNIFIED STORAGE SYSTEMS

Abstract

This white paper discusses strategies for migrating data to EMC[®] VNX[™] series storage systems. Considerations include mixing 4 Gb/s and 8 Gb/s connectivity, tools for moving data, and planning for and minimizing downtime.

May 2011



Copyright © 2011 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate 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.

VMware is a registered trademark of VMware, Inc. All other trademarks used herein are the property of their respective owners.

Part number h8248



Table of Contents

Executive summary
Audience4
Introduction
EMC VNX unified storage systems
Connectivity
Using Fibre Channel as a data path7
Using a host-based data path8
Source storage9
Data transfer throughput10
Downtime10
Migration tool selection
Storage-system-based migration tools (Fibre Channel data path)12
SAN Copy12
Open Replicator13
MirrorView14
Host-based migration tools15
PowerPath Migration Enabler15
VERITAS Volume Manager for Solaris16
AIX LVM mirroring16
Tool properties for consideration17
Conclusion
References



Executive summary

The EMC[®] VNX[™] series of storage systems is a member of the VNX family and is EMC's next-generation midtier unified storage platform. It is designed for midtier-toenterprise storage environments that require advanced features, flexibility, and configurability. The EMC VNX series of midrange storage systems features major architectural innovations such as the new SAS back end. The point-to-point architecture of the 6 Gb SAS back-end infrastructure offers high availability and enhanced performance. The architecture of the front end enables the VNX unified series to deliver full 8 Gb/s capabilities throughout the entire system. Since SANs are currently based on different technologies, EMC unified storage systems are designed to support equipment with mixed speeds on the same SAN fabric. This is done with expanded UltraFlex[™] I/O, which supports FcoE, FC, ISCSi, CIFS, NFS [Parallel-NFS (pNFS), Multi-Path File System (MPFS)] all on one single path.

Migration of existing data to a VNX series storage system may be done for a number of reasons, including:

- Consolidating existing storage on a common platform
- Taking advantage of the increased performance of the VNX architecture
- Leveraging the newly added features such as FAST VP, FAST Cache and so on
- Reallocating data from one site to another site

This paper explores the many options that are available for migrating existing data to VNX series storage systems. EMC VNX unified storage systems include specific tools for transferring data, such as SAN Copy[™] and MirrorView[™] software. In addition, there are many other tools for migrating data that are described in this paper. Ultimately, the best migration tool is the tool that can transfer data with minimal downtime.

The purpose of this paper is to help determine which of the many available migration tools is best for the situation.

Audience

This white paper is intended for EMC employees, partners, and customers who are planning to migrate existing data to EMC VNX series storage systems. This paper is targeted to an audience that has some level of familiarity with EMC CLARiiON[®] CX UltraScale[™] series storage systems as well as EMC VNX series storage systems.

Introduction

This white paper outlines the planning considerations for a migration to the EMC VNX series unified storage system, and will help you select the appropriate tools to complete the migration. Considerations include:

• Connectivity



The EMC VNX series provides 8 Gb/s SAN connectivity. There are several ways the VNX series can be connected to existing storage to allow data transfer between the storage systems.

Downtime

Most migration solutions require some application downtime. Some migration solutions require downtime for the entire duration of the data transfer, while other solutions require only enough downtime for a host reboot. The downtime required by different migration tools is discussed in this paper.

Migration tools

Storage-system-based and host-based tools are available to transfer data. This paper contains descriptions of several tools to help you choose the best tool for the application.

The "Connectivity" and "Downtime" sections outline factors to consider before beginning the migration. The "Migration tool selection" section describes several migration tools and their requirements for connectivity and downtime.

This paper is not a how-to guide; separate user guides are available for each of the tools discussed in this paper. The purpose of this paper is to help you choose the best migration tool for your application.

A typical customer environment consists of many hosts, storage systems, and LUNs. For simplicity, the migration examples in this paper are presented as a single source to a single destination.

EMC VNX unified storage systems

The VNX series is a member of the VNX family, which combines all the benefits of Celerra[®] and CLARiiON storage systems into a single product line.

Figure 1 shows the new generation of EMC unified storage.





Figure 1. EMC VNX family of unified storage systems

VNX is a "one box fits all" platform for many reasons:

- VNX offers support for any type of storage protocol (iSCSI, Fibre Channel, or Oracle DNFS) that customers use in their test, development, and production Oracle database environments.
- VNX provides easy migration of Oracle environments between Fibre Channel (FC) and IP protocols. With EMC and Oracle replication technologies, users can execute heterogeneous replication in multiprotocol environments to benefit specific workloads, such as putting disaster recovery storage onto a unified platform.
- VNX enables a fast recovery option for Oracle workloads. Replication Manager is a wizard-driven interface that interfaces with Oracle and provides nearly instantaneous copies of data. And with a "gold copy" of space-efficient snapshots and clones, customers can recover an entire database in minutes. For remote migration most reliable EMC technologies such as MirrorView and SAN Copy can be leveraged. Different EMC VNX replication methodologies discussed in EMC Unified Storage Solutions: Oracle Database 11gR2 with EMC VNX and Storage Replication Consistency for use with Oracle extend to all models in the EMC VNX family that support SnapView and MirrorView and can be automated through shell scripts and batch files using EMC Navisphere[®] CLI commands.
- FAST VP analyzes disk I/O patterns, automatically promotes sub-LUN "hot" spots to a higher performance drive, and demotes cold data to the high-capacity drives.
- The EMC FAST Cache feature provides a further performance boost to Oracle workloads by increasing read and write cache, resulting in turbo-charging the existing architecture.
- VNX also supports the concept of pools introduced in previous releases of EMC FLARE[®] (now the VNX Operating Environment) that allow customers to readily



transform their IT functions by fully harnessing the advantages of an enterprise private cloud.

The state-of-the-art I/O interconnection based on native PCI Express technology within the UltraFlex architecture delivers high bandwidth, low latency characteristics that enable the EMC VNX series to deliver full 8 Gb/s capabilities throughout the entire system. The following are EMC VNX series unified storage-system features:

- It supports both direct-attach and SAN environments.
- It uses hot-swappable storage processors with up to 8 GB of memory per SP.
- It includes 8 Gb/s front-end ports.
- It includes 6 Gb/s SAS back-end ports.
- It supports 4 Gb/s and 8 Gb/s disk drive interfaces.
- It supports RAID level 0, 1, 1/0, 3, 5, 6, and individual disks, along with global hot sparing.
- It has nondisruptive upgrade (NDU) capability.
- It supports Windows, Solaris, Linux, AIX, HP-UX, and VMware host attaches.

Connectivity

To move data from the source storage to the VNX storage system, a data path between the source storage device and the VNX storage system must exist. The data path can be a Fibre Channel connection between the two storage devices, or a connection at the host level, such as an IP connection. The migration tool chosen will depend partly on what data paths are available.

Using Fibre Channel as a data path

Tools such as SAN Copy, Open Replicator, and MirrorView transfer data between storage systems over Fibre Channel. They require a fibre connection between the front-end ports of the storage systems (SP ports on the EMC VNX unified). The connection can be either a direct-attach connection or a SAN connection. In our architecture we used a Brocade 5300 switch.

When mixing different types of hardware on a SAN, SAN components can be running at 1 Gb/s, 2 Gb/s, 4 Gb/s, or 8 Gb/s speeds. The EMC VNX unified series SP ports auto-negotiate to the maximum possible speed for that configuration. For example, if a VNX storage system is connected to a 8 Gb/s switch, its front-end ports will auto-negotiate to 8 Gb/s, even if other components connected to that switch are connected at 4 Gb/s.

Figure 2 is a simplified example in which the CX4-480 storage systems connect to the SAN at their maximum speed of 4 Gb/s, while the VNX still maintains full 8 Gb/s connectivity on all ports. This allows for other 4 Gb/s SAN hardware (such as servers)



to transfer data to and from the VNX at full 8 Gb/s speeds. EMC recommends that the VNX SP ports be set to auto-negotiate to minimize the potential for speed conflicts.



Figure 2. Mixing Fibre Channel speeds

Using a host-based data path

Other tools transfer data at the host level using the host-based data path as shown in Figure 3. These tools can be used to migrate when a single host has access to both source and destination storage. The host reads data from the source storage and writes it to the destination storage. With the right migration tool, this data transfer is performed without application downtime and the only potential downtime is when



the source-to-destination relationship is established or severed.



CX4-480

Figure 3. Using a host-based data path

Source storage

The storage currently storing the data to be migrated is known as the "source storage." In some cases, the location of the source data determines the data path, as shown by the following four examples:

• Source data is on internal server disks

In this example, the source storage is not SAN-connected and the only usable data path is host-based.

• Source data is on non-EMC (third-party) SAN storage

If the third-party storage is qualified¹ with SAN Copy, then the SAN is a valid data path.

A host-based data path is also valid.

• Source data is on a previous generation CLARiiON storage system

When the source storage is on a CLARiiON storage system, tools such as SAN Copy and MirrorView can be used to migrate data. Be sure to check SAN Copy or MirrorView support for the particular storage system.

A host-based data path is also valid.

• Data is on an EMC Symmetrix[®] storage system

A Fibre Channel data path is possible with either SAN Copy or Open Replicator.

¹See *E-Lab Navigator* on <u>EMC Powerlink</u>[™] for the latest interoperability matrices. Choose the 'PDFs and Guides' tab, then choose the 'SAN Copy' PDF.



A host-based data path is also valid.

Data transfer throughput

The process of moving data from the source to the destination includes preparing to transfer the data, transferring the data, and reconfiguring the hosts to access the data in its new location.

Every migration is affected by how fast data can be moved from the source storage to the destination storage (the storage device to which the data is going to be migrated). In this paper, the destination storage is always a VNX series storage system. Transfer throughput can be limited by how quickly data can be read from the source or written to the destination. It can also be limited by the transfer medium between the source and destination.

If possible, test migration in advance to get an estimated throughput. Some migration tools require application downtime for the duration of the migration. In that case, it is beneficial to get the highest throughput possible to minimize downtime. Other migration tools can copy data in the background while the application is online. In this case, the duration of the transfer may not have a significant impact on downtime.

Downtime

Downtime is the amount of time during the migration that end users will not have access to the data. Almost every migration requires application downtime at some point. Sometimes downtime is required when setting up for the migration. Things that can introduce downtime during migration preparation are:

• Upgrading server hardware

Servers must be powered down to install 8 Gb/s HBAs. This can be done either before or after other SAN components are upgraded to 8 Gb/s, depending on interoperability of the equipment². EMC E-Lab[™] has qualified several 8 Gb/s HBAs for connectivity with previous generation CLARiiON storage systems, so it is possible to take full advantage of 8 Gb/s connectivity to the EMC VNX unified series while maintaining compatibility with existing CLARiiON CX4 storage systems.

• Upgrading SAN hardware

Introducing 8 Gb/s switches into a fabric may require a disruption in connectivity while rewiring or rezoning the fabric.

Software upgrades

Some migration solutions require software installation that needs a host reboot before the software becomes active.

² See *E-Lab Navigator* on <u>EMC Powerlink</u>[®] for the latest interoperability matrices.



Sometimes downtime is required during the transfer of data. Some factors that cause downtime during the migration are:

• Application quiesce (to momentarily pend I/O to a device) requirements

Depending on the migration solution, the application may need to be quiesced for varying amounts of time. Sometimes the application must be quiesced for the duration of the transfer, sometimes momentarily, and sometimes not at all. Application quiesce requirements will be discussed in the next section with respect to each tool.

• Mapping to a new destination volume

The destination volume will have a different identity than the source volume. This will require:

- Zoning The host must be zoned to the EMC VNX unified series storage system.
- LUN masking The destination LUN must be placed in a storage group for the host.
- Host LVM mappings The destination must be given a drive letter or mount point on the host.
- Application data store mappings If the destination mount point is different than that of the source, the application's configuration must be updated to point to the new location of its data.
- Raw partitions to create ASM disk group If using the raw partition to store the data in case of the source data, then once the migration is completed, raw partitions must be mapped to exact corresponding destination devices

For example, if the source, emcpowera1, is mapped to raw device raw1 and emcpowerb1 is the device on the destination storage that corresponds to emcpowera1, then emcpowerb1 should be mapped to raw1 in order to mount the disk group using those raw devices

 Oracle ASM driver - If an Oracle asm driver is used to mark the devices, then once the LUNs are masked check if all the LUNs are available and then mount the diskgroups using the same name as earlier

Some of these steps may be completed prior to the migration to minimize downtime while the rest can be done once the migration is completed.

Migration tool selection

Many tools are available to handle the task of transferring data. This paper breaks these tools into two major categories, storage-system-based migration tools that use a Fibre Channel data path, and host-based migration tools that use a host-based data path.



Regardless of the migration tool selected, EMC always recommends a full backup prior to any migration.

Storage-system-based migration tools (Fibre Channel data path)

Storage-system-based migration tools allow data to be migrated directly across a Fibre Channel connection between storage systems. Some advantages of these migration tools include:

- The host does not have to do any of the work; all data transfer is performed by the storage systems.
- Since there is often no software to install on the host, host reboots can potentially be reduced.
- The same migration tool can be used across heterogeneous operating systems.

Some possible disadvantages are:

- Depending on the configuration, the source LUN may have to be quiesced for the duration of the migration. This results in a longer application downtime.
- The host will have to remap to the new destination LUN, which will require the application to be taken down at some point.
- The unit of migration is a whole LUN. This may be undesirable if only some of the data on the LUN needs to be migrated.

SAN Copy

SAN Copy is a storage-system-based data mobility tool that runs on an EMC VNX unified storage system. It moves LUNs across Fibre Channel directly between storage systems. It has two modes of transfer – full transfers and incremental transfers. Full transfers require the source and destination LUNs to be unmounted for the duration of the transfer in order to get coherent data at the destination. Incremental transfers can be performed with the source LUN online.

SAN Copy also has two directions of transfer – a push or a pull. A pull occurs when the VNX storage system running SAN Copy reads data from another storage system and writes it to one of its own LUNs. A pull is always a full transfer. A push occurs when a CLARiiON storage system running SAN Copy reads data from one of its own LUNs and writes it to the EMC VNX storage system. A push can be full or incremental.

Considering these modes of operation for SAN Copy, an incremental push is best for migration purposes because it requires downtime only to perform the last incremental synchronization and to remap the host to the destination LUN. Pulls, on the other hand, require more downtime because the source storage must be offline for the duration of the data transfer. In either case, the destination should not be mounted during the data transfer.



SAN Copy is ideal as a migration tool when the source storage is a CLARiiON system that can run SAN Copy. This allows incremental pushes to be used, which minimizes downtime for a SAN Copy migration solution.

The basic steps for a migration using SAN Copy are:

- 1. Create an incremental SAN Copy session for the source and destination LUNs.
- 2. Start the session with the following Navisphere CLI command (will be supported in Unisphere GUI for the future releases of the VNX OE):

sancopy -start -name sessionName -copywholelun -nomark all

This is an efficient method for performing an incremental transfer because it skips copy-on-first-write operations that are usually performed on the source. This results in a faster transfer with less performance impact on the application but requires a follow-up incremental transfer to bring the destination storage to a consistent state.

- 3. (Optional test) For a moment, quiesce the application while starting an incremental transfer. When the transfer is complete, attempt to restart the application from a snapshot of the destination. If this step is omitted and the application fails to restart at the destination, another full transfer may be necessary to clean up any changes to the destination. A snapshot prevents the destination from being changed inadvertently.
- 4. Shut down the application and unmount the source LUN. At this point minimal differences should exist between the source and destination, so the time before the destination can be mounted and the application brought back online is minimal. To determine the percentage of the LUN that needs to be transferred, run the following Navisphere CLI command and compare **Number of Blocks to Copy** with **Size of source LUN**:

sancopy -- info -- name sessionName -- all

5. Start an incremental transfer with the following command:

sancopy -start -name sessionName

6. When the transfer is finished, mount the destination volume and restart the application.

Open Replicator

Open Replicator is a data movement tool that runs on a Symmetrix storage system. Within the context of this paper, it is most useful as a migration tool when the source storage is on Symmetrix because differential copying (similar to SAN Copy incremental copying) can be used to minimize application downtime.

Differential copying transfers only changed blocks after a full initial transfer. To start a full transfer where subsequent transfers will be differential, use the **differential** switch. For example:

symrcopy create -copy -name sessionLabel -push -hot -file pairs -differential

symrcopy activate -file pairs



where *pairs* is a file containing source-to-destination volume mappings. Subsequent incremental transfers should be run using the recreate command:

symrcopy recreate -name sessionLabel -file pairs

symrcopy activate -file pairs

Using SYMCLI with EMC Open Replicator for Symmetrix, available on <u>EMC Powerlink</u>, provides more detailed information about how to control Open Replicator from the command line.

MirrorView

MirrorView/Synchronous and MirrorView/Asynchronous are migration tools that can be used when the source is a CLARiiON storage system and destination is an EMC VNX storage system. Both products synchronize the destination LUN with the source LUN without any application downtime and require downtime only while remapping the host from source to destination. MirrorView/Asynchronous requires slightly more downtime because a final incremental synchronization must be performed when the application is being switched over to the destination. Once the destination is synchronized with the source, an application failover can be performed to move the application over to the destination.

When using MirrorView as a migration tool, it is important to remember that writes to the destination LUN following a promote operation will be mirrored back to the source volume. If a problem occurs during application restart at the destination, potentially unwanted changes may be mirrored back to the source, corrupting the original copy of the data. Do not make use of the MirrorView promote operation during a migration and follow the instructions below to avoid this potential pitfall.

MirrorView is a migration option when the source is on a previous generation CLARiiON.³ The process is similar for both MirrorView/Synchronous and MirrorView/Asynchronous. The steps for this procedure are described below:

- 1. Enable the MirrorView paths between the source and destination storage using the Navisphere CLI commands (GUI support for Mirror View will be incorporated in future VNX Operating Environment releases)
- Create a MirrorView mirror for the source and destination LUNs. For MirrorView/Asynchronous set update type to manual update. Allow for the initial synchronization to complete (this may take a significant amount of time depending on how much data needs to be transferred).
- 3. Shut down the application and unmount the source LUN.
- 4. (For MirrorView/Asynchronous only) Perform a synchronization to synchronize the destination LUN with the source LUN. Only incremental changes will be transferred.

³ MirrorView is qualified for use between systems running the same major VNX OE release, or within one major VNX OE release. Check *E-Lab Navigator* on <u>EMC Powerlink</u> for the latest hardware interoperability matrices.



- 5. Wait for the mirror to reach the **synchronized** state.
- 6. Fracture the mirror.
- 7. Attempt to restart the application from a snapshot of the destination. This step is optional; however, if when skipping this step, the application fails to restart at the destination, a full resynchronization may be necessary to clean up any changes to the destination. A snapshot prevents the destination from being changed inadvertently.
- 8. If the test restart succeeds, remove the secondary image from the mirror and restart the application directly from the destination LUN. Otherwise, restart the application from the source, synchronize the mirror, and return to step 3.

Host-based migration tools

Host-based migration tools transfer data from one location to another at the host level. Some advantages of host-based migration tools are:

- Host-based migration tools can easily move data between different types of storage devices.
- The unit of migration is a file rather than a whole LUN, so individual files can be migrated.
- Using tools that are tightly integrated with the application or file system may result in less downtime than a storage-system-based migration solution.

Some possible disadvantages are:

- The tool has to be installed on each host that is migrating data. If not already installed, it may require a reboot to become active.
- A different migration tool may be required for each operating system.
- Many host-based tools require both source and destination storage to be visible to the same host. If the application is migrated to a new server as well as new storage, the migration will have to be performed in two steps, one to change the storage and another to change the server.

PowerPath Migration Enabler

The host-based EMC PowerPath[®] Migration Enabler (PPME) is integrated with other technologies to eliminate application downtime during data migrations or virtualization implementations. It works with EMC Open Replicator for Symmetrix to synchronize storage systems during data migrations, or with EMC TimeFinder/Clone to move data across tiers, with minimal impact to host resources. Additionally, PowerPath Migration Enabler host copy eliminates downtime by enabling nondisruptive data migrations through the host in heterogeneous environments, without the need for an underlying migration technology.

PPME host copy supports multiple migration scenarios such as standard to standard devices, virtually provisioned devices encrypted devices, pseudo and native devices. The detailed description is out of scope of this paper; see the "References" section.



Every I/O to the host has to pass through PowerPath since it resides below logical volume managers, file systems and applications but above the SCSI drivers. This allows PowerPath to work with the storage system to provide efficient I/O management.

In order to use the PPME, first install PowerPath and then install PowerPath Migration Enabler . Once the installation is done, zone and mask the target LUN to the host. The two primary components in Migration Enabler are the user space and the kernel space modules. PPME executes the CLI commands in the user space and so migration control is given to that component, but the movement of data occurs in the kernel space. After identifying the source and the target LUNs, use the following commands to proceed with the migration:

1. Set up migration with:

powermig setup

2. Set up copying and cloning with:

powermig sync

3. Check the status with:

powermig query

4. Test the results after copying the data with:

powermig selectTarget

5. Finally commit the migration with:

powermig commit

6. Pause or resume the migration as required with:

powermig stop/resume

7. Once the results indicate that the migration is complete, clean up the migration and remove the source device.

VERITAS Volume Manager for Solaris

VERITAS Volume Manager provides storage management capabilities within a Sun Solaris environment. It is supported with EMC storage through the EMC foundation suite. It is an ideal candidate as a migration tool when it is already being used to manage the source storage devices and the destination storage can be connected to the same volume group. One limitation of using VERITAS Volume Manager is that RAID 5 volumes cannot be mirrored and are therefore not a candidate for this migration tool. Otherwise, this tool provides high levels of data availability and data integrity during the migration.

AIX LVM mirroring

Logical Volume Manager (LVM) is a subsystem within AIX that manages the disk environment and provides easy-to-use, online storage management. The AIX LVM is



managed either through the AIX management tool (SMIT) or directly from the command line. It is an ideal candidate for a migration tool when it is already being used to manage the source storage devices and the destination storage can be connected to the same volume group.

Tool properties for consideration

When selecting a tool to use for migration certain properties will influence which migration tool is right for the given project, including:

• Setup complexity

The amount of effort it takes to implement the tool is a consideration, especially if the tool is only going to be used temporarily. Storage-system-based migration tools may require changes to SAN layout and zoning, while host-based tools may require an installation and host reboots to configure the tool properly.

• Downtime required during data transfer

Some migration tools require the source volume to be offline for the duration of data transfer (SAN Copy pulls, copy commands); others can transfer data while the source is active (SAN Copy incremental push, Open Replicator differential copying and LVM mirroring).

• Downtime required to map to the destination

When the data movement is complete, the source must be removed from service and the application must be restarted on the destination. Host-based LVM mirroring minimizes this because the application is already running on both the source and destination when the data transfer is complete.

Scaling

It is important to verify that the migration tool can handle the number of volumes that need to be replicated. For example, SAN Copy can transfer up to 16 LUNs (if from CX4-480,CX4-960 to VNX) simultaneously, but it can have hundreds of LUNs queued to be transferred. MirrorView/A and MirrorView/S can have up to 512 mirrors. With either MirrorView/A or MirrorView/S software, up to 40 mirrors can be synchronizing at the same time.



Conclusion

EMC provides both storage-system-based and host-based migration tools, as well as a range of services to assist with the process of migrating data to these new storage systems. Also EMC supports several non-EMC data migration tools for data migration. Downtime is usually the most important consideration for a migration solution. The information in this paper provides useful guidelines for selecting the proper migration tool to minimize application downtime.

References

The following white paper can be found on EMC.com:

- EMC PowerPath Migration Enabler Host Copy
- <u>EMC Unified Storage Solutions: Oracle Database 11gR2 with EMC VNX and</u> <u>Storage Replication Consistency</u>
- Leveraging EMC Unified Storage System Dynamic LUNs for Data Warehouse
 Deployments
- Migrating Data from an EMC CLARiiON Array to a VNX Platform using SAN Copy

