

Data Migration Techniques for VMware vSphere

A Detailed Review

EMC Information Infrastructure Solutions

Abstract

This white paper profiles and compares various methods of data migration in a virtualized environment. In-array, cross-array, and host-based methods are examined. While the primary focus of this paper is on the VMware vSphere 4.1 infrastructure, it also assesses replication options, storage virtualization, and the tools that can assist administrators in migrating not just storage, but also the applications and services utilizing that storage. Information about the most appropriate replication strategy for each different scenario is provided.

November 2010

Copyright © 2010 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

All other trademarks used herein are the property of their respective owners.

Part number: H8063

Table of Contents

Executive summary	6
Business case	6
Solution overview	6
Key results.....	6
Introduction	7
Introduction to this white paper	7
Purpose.....	7
Scope	7
Audience	8
Terminology.....	8
Technology overview	9
Introduction.....	9
EMC CLARiiON CX4-480	9
EMC Unisphere	9
EMC VPLEX Metro.....	9
EMC VPLEX Local	10
EMC RecoverPoint.....	10
EMC SAN Copy.....	10
CLARiiON LUN Migrator	10
VMware vSphere.....	10
VMware vCenter SRM	10
VMware vCenter Converter.....	11
Configuration	12
Overview	12
Physical environment	12
Hardware resources	13
Software resources	13
Data migration with VMware vCenter Storage vMotion	14
Overview	14
Administering the virtual infrastructure.....	14
Migration options.....	15
Destination disk types in Storage vMotion	15
Using VMware Storage vMotion.....	16
Storage vMotion requirements and limitations.....	19
Storage hardware acceleration	19
Comparing Storage vMotion with and without VAAI enabled	21
Conclusion.....	24
Data migration with VMware vCenter Converter.....	25
VMware vCenter Converter.....	25

Hot cloning with synchronization and switch features.....	25
Configuring a hot cloning operation	26
Performance impact of hot cloning.....	32
Cold cloning.....	33
Data migration with EMC VPLEX Metro	34
Introduction to the EMC VPLEX family	34
Data migration with VPLEX.....	35
Configuration of VPLEX Metro	36
VPLEX virtual volumes.....	37
Introducing VPLEX into an existing environment.....	37
Migration scenario	38
Conclusion.....	41
Data migration with CLARiiON SAN Copy	42
Introduction to EMC CLARiiON SAN Copy	42
Environment topology	42
Uses for SAN Copy	43
Migrations available with SAN Copy	43
Benefits of SAN Copy	43
SAN Copy full and incremental sessions	43
SAN Copy in a VMware environment.....	44
Requirements and considerations.....	44
Migration duration and performance	44
CLARiiON SAN Copy setup overview.....	46
Migrating to third-party storage arrays	47
Migration scenario	48
Conclusion.....	54
Data migration with CLARiiON LUN Migrator	55
CLARiiON virtual LUN Migrator.....	55
Migration duration and performance	56
CLARiiON LUN migration procedure	56
Scenario: Migrating applications with CLARiiON LUN Migrator	58
Environment topology	59
Array performance during the migration	60
Application performance during the migration	61
Conclusion.....	62
Data migration with VMware vCenter Site Recovery Manager.....	63
Test environment.....	63
Requirements for using VMware vCenter SRM as a migration tool	64
Using VMware vCenter SRM as a migration tool.....	65
Installing and configuring VMware vCenter SRM failovers.....	65
Impact on production performance	67
Using RecoverPoint SRA for local array or LUN migration.....	69
Summary of data migration techniques	70

Comparing the data migration techniques	70
Scenario 1: Migration of individual virtual disks	70
Scenario 2: Migration of datastore contents to another LUN on same array.....	71
Scenario 3: Migration of entire contents of a datastore to a LUN on another array	71
Scenario 4: Migration of the entire virtual infrastructure to another geographical location.....	72
Scenario 5: Migration from physical to virtual or inter-hypervisor	73
Conclusion	74
Summary	74
Findings.....	74
Next steps	74
References.....	75
White papers	75
Other documentation.....	75

Executive summary

Business case To meet the business challenges presented by today's on-demand 24x7 world, data must be highly available—in the right place, at the right time, and at the right cost to the enterprise. IT organizations are increasingly being tasked with increasing flexibility and agility within the enterprise and at the center of those capabilities is data migration. VMware vSphere increases flexibility at the server level. For flexibility to be realized across the enterprise, the data and storage must mirror that flexibility. Data migration must occur seamlessly and without impacting applications or end users.

Solution overview A variety of techniques and tools are available to customers when migrating virtual data centers, each with its own advantages and disadvantages. This white paper investigates several of the methods commonly used for virtual data center migrations. It mainly examines the VMware vSphere 4.1 infrastructure but also assesses replication options, storage virtualization, and the tools that can assist administrators in migrating not just storage but also the applications and services using that storage.

The main focus areas are:

- Migrating virtual environments using native VMware and EMC tools and functionality, such as VMware vCenter Converter, VMware Storage vMotion, EMC® CLARiiON® LUN Migrator, and EMC CLARiiON SAN Copy™
 - Disaster recovery (DR) with VMware vCenter Site Recovery Manager (SRM) as a migration tool for coordinating, testing, and executing a data center migration
 - EMC VPLEX™ Metro, which provides campus-based migration within and/or between data centers across distances of up to 100 km
 - Replication options relating to distance, protocol, and scheduled downtime
-

Key results This comparative study highlights several approaches to data migration for VMware vSphere and provides valuable insight for customers planning a data center migration of their virtual information infrastructure environments. Customers are able to fully realize data migration flexibility at the server and storage layers of the environment.

Introduction

Introduction to this white paper

In this white paper, valid scenarios are suggested and tested for a number of data migration methods using VMware and EMC tools and functionality. This paper documents the impact to these applications during data migration. This white paper includes the following sections:

Topic	See Page
Technology overview	9
Configuration	12
Data migration with VMware vCenter Storage vMotion	14
Data migration with VMware vCenter Converter	25
Data migration with EMC VPLEX Metro	34
Data migration with CLARiiON SAN Copy	42
Data migration with CLARiiON LUN Migrator	55
Data migration with VMware vCenter Site Recovery Manager	63
Summary of data migration techniques	70
Conclusion	74
References	75

Purpose

The purpose of this white paper is to profile and compare various methods of data migration in a virtualized environment. It examines in-array, cross-array, and host-based methods and provides information about the most appropriate replication strategy for each different scenario.

Scope

The objectives of this white paper are to examine suitable scenarios and configurations for:

- VMware Storage vMotion
- VMware vCenter Converter
- EMC VPLEX Metro
- EMC SAN Copy
- EMC CLARiiON LUN Migrator
- VMware vCenter SRM with EMC RecoverPoint

Included are representative virtualized Exchange and SQL implementations. Actual customer implementations may vary from the parameters shown in this paper, based on testing results.

Audience

This white paper is intended for:

- Field personnel who are tasked with deploying similar solutions
- Customers, including IT planners, storage architects, and those deploying similar solutions
- EMC staff and partners, for guidance and the development of proposals

This paper assumes that you are familiar with:

- VMware and vSphere 4 technology
 - EMC VPLEX Metro and RecoverPoint
-

Terminology

This section defines terms used in this document.

Term	Definition
DR	Disaster recovery.
SP	Storage processor on a CLARiiON storage system. On a CLARiiON storage system, a circuit board with memory modules and control logic that manages the storage-system I/O between the host's Fibre Channel (FC) adapters and the disk modules.
Unisphere	EMC Unisphere™ software provides the next generation of storage management and presents a single, integrated, and simple web interface for unified storage arrays, as well as standalone CLARiiON and Celerra® storage systems.
Virtual machine	A software implementation of a machine that executes programs like a physical machine.
VMDK	Virtual Machine Disk format. A VMDK file stores the contents of a virtual machine's hard disk drive. The file can be accessed in the same way as a physical hard disk.
VMware vCenter SRM	VMware vCenter Site Recovery Manager.

Technology overview

Introduction This section briefly describes the key technologies deployed in the test environment.

EMC CLARiiON CX4-480 The EMC CLARiiON CX4 series delivers industry-leading innovation in midrange storage with the fourth-generation CLARiiON CX storage platform. The unique combination of flexible, scalable hardware design and advanced software capabilities enables the CLARiiON CX4 series systems to meet the growing and diverse needs of today's midsize and large enterprises. Through innovative technologies like Flash drives, UltraFlex™ technology, and CLARiiON Virtual Provisioning™, customers can:

- Decrease costs and energy use
- Optimize availability and virtualization

CLARiiON CX4-480 is a versatile and cost-effective solution for organizations seeking an alternative to server-based storage. It delivers performance, scalability, and advanced data management features in one, easy-to-use storage solution.

EMC Unisphere EMC Unisphere provides a flexible, integrated experience for managing existing CLARiiON storage systems and next-generation EMC unified storage offerings in a single screen. This new approach to midtier storage management fosters simplicity, flexibility, and automation. Unisphere's unprecedented ease of use is reflected in intuitive task-based controls, customizable dashboards, and single-click access to realtime support tools and online customer communities.

EMC VPLEX Metro EMC VPLEX Metro enables disparate storage arrays at two separate locations to appear as a single, shared array to application hosts, enabling customers to easily migrate and plan the relocation of application servers and data, whether physical or virtual, within and/or between data centers across distances of up to 100 km.

VPLEX Metro enables companies to ensure effective information distribution by sharing and pooling storage resources across multiple hosts over synchronous distances.

VPLEX Metro empowers companies with new ways to manage their virtual environment over synchronous distances so they can:

- Transparently share and balance resources across physical data centers
 - Ensure instant, realtime data access for remote users
 - Increase protection to reduce unplanned application outages
-

EMC VPLEX Local EMC VPLEX Local provides seamless data mobility and lets organizations manage multiple heterogeneous arrays from a single interface within a data center. VPLEX Local provides a next-generation architecture that enables customers to increase availability and improve utilization across multiple arrays.

EMC RecoverPoint EMC RecoverPoint supports cost-effective, continuous data protection and continuous remote replication for on-demand protection and recovery to any point in time. RecoverPoint's advanced capabilities include policy-based management, application integration, and bandwidth reduction.

RecoverPoint provides a single, unified solution to protect and/or replicate data across heterogeneous storage. With RecoverPoint, organizations can simplify management and reduce costs, recover data at a local or remote site to any point in time, and ensure continuous replication to a remote site without impacting performance.

EMC SAN Copy EMC CLARiiON SAN Copy is a storage-system-based application that is available as an optional package. SAN Copy is designed as a multipurpose replication product for data mobility, migrations, content distribution, and disaster recovery. SAN Copy enables the storage system to copy data at a block level directly across the SAN, from one storage system to another, or within a single CLARiiON system. While the software runs on the CLARiiON storage system, it can copy data from, and send data to, other supported storage systems on the SAN.

CLARiiON LUN Migrator CLARiiON LUN Migrator is a feature that moves data, without disruption to host applications, from a source LUN to a destination LUN of the same or larger size, and with requisite characteristics within a single storage system. LUN migration leverages EMC FLARE[®], CLARiiON's existing operating system, for data integrity and RAID protection features. The functions are integrated into EMC Unisphere and CLI packages. The driver that facilitates the LUN migration operations is packaged with the FLARE operating environment.

VMware vSphere VMware vSphere is the industry's most complete, scalable, and powerful virtualization platform, delivering the infrastructure and application services that organizations need to transform their information technology and deliver IT as a service. VMware vSphere provides unparalleled agility, control, and efficiency while fully preserving customer choice.

VMware vCenter SRM VMware vCenter SRM is a DR management and automation solution for VMware virtual infrastructures. VMware vCenter SRM accelerates recovery by orchestrating and automating the recovery process and simplifying management of disaster recovery plans.

**VMware
vCenter
Converter**

VMware vCenter Converter is a highly robust and scalable enterprise-class migration tool that automates the process of creating VMware virtual machines from physical machines, other virtual machine formats, and third-party image formats. Through an intuitive wizard-driven interface and a centralized management console, VMware vCenter Converter can quickly and reliably convert multiple local and remote physical machines, without any disruptions or downtime.

Configuration

Overview

The following section identifies and briefly describes the technology and components used in the test environment.

Physical environment

The following diagram provides an example of the overall physical architecture of the environment.




Figure 1: Environment overview

Hardware resources

The hardware used to validate the solution is listed in Table 1.

Table 1: Hardware requirements

Equipment	Quantity	Configuration
Servers	3	<ul style="list-style-type: none">• 16-core• 128 GB RAM• 4 NICS
Ethernet switch	1	Ethernet 1 Gb switch 48-port
Storage	2	EMC CLARiiON CX-480 Drive count: 10 x 300 GB FC
Storage	2	VPLEX Metro
EMC RecoverPoint appliances	4	Gen 4 appliances
FC switch	2	Departmental switches

Software resources

The software used to validate the solution is listed in Table 2.

Table 2: Software requirements

Software	Version
EMC CLARiiON FLARE	FLARE 30
EMC PowerPath®/VE	5.4.1
EMC RecoverPoint	3.3
EMC SAN Copy	FLARE 30
EMC Unisphere	1.0
VMware vCenter	4.1
VMware vCenter Converter Standalone	4.3
VMware vCenter Site Recovery Manager	4.1
VMware vSphere	4.1
Microsoft Exchange 2010	RTM (Build 14.00.0639.021)
Microsoft SQL Server 2005	Enterprise Edition
Microsoft Windows 2008	Enterprise Edition

Data migration with VMware vCenter Storage vMotion

Overview

With VMware vCenter Storage vMotion, a virtual machine and its disk files can be migrated from one datastore to another while the virtual machine is running. These datastores can be on the same storage array, or they can be on separate storage arrays. Storage vMotion is supported for use with FC, network file system (NFS), and Internet small computer system interface (iSCSI) storage protocols.

Administering the virtual infrastructure

Storage vMotion can be used in several ways to administer the virtual infrastructure:

- Storage maintenance and reconfiguration—Storage vMotion can be used to move virtual machines off a storage device to enable maintenance or reconfiguration of the storage device without virtual machine downtime.
- Redistributing storage load—Storage vMotion can be used to manually redistribute virtual machines or virtual disks to different storage volumes to balance capacity or improve performance.
- Upgrading VMware ESX/ESXi without virtual machine downtime—During an upgrade from ESX server 2.x to ESX/ESXi 3.5 or later, running virtual machines can be migrated from a VMFS2 datastore to a VMFS3 datastore. The VMFS2 datastore can be upgraded without any impact on the virtual machines. Storage vMotion can then be used to migrate the virtual machines back to the original datastore without any virtual machine downtime.

From an infrastructure perspective, the main requirement is that both the source and target datastores are accessible to the ESX/ESXi host on which the virtual machine is hosted as shown in Figure 2. The virtual machine does not change execution host during a migration with Storage vMotion.




Figure 2: Migration with Storage vMotion

Migration options

Depending on the running state of the virtual machine, there is a slight difference in the options available to the user. A powered-off virtual machine provides the full range of migration options that can occur simultaneously, whereas a powered-on virtual machine is restricted to migrating either the resources or the data in the same job.

These options are detailed in Table 3.

Table 3: Migration options

Powered-off or suspended virtual machine	
Change host	Move the virtual machine to another ESX/ESXi host
Change datastore	Move the virtual machine's configuration file and virtual disks
Change both host and datastore	Move the virtual machine to another ESX/ESXi host and move its configuration file and virtual disks
Powered-on virtual machine	
Change host	Move the virtual machine to another ESX/ESXi host
Change datastore	Move the virtual machine's configuration file and virtual disks

To execute a VMware Storage vMotion operation, select the **Migrate** option from the context menu of the virtual machine to be migrated. The same menu option is selected when executing a VMware vMotion operation. Table 3 details the migration options presented.

Note The **Change host** option is a vMotion operation only, not a Storage vMotion operation. You cannot perform vMotion and Storage vMotion simultaneously on a running virtual machine. This testing ran on powered-on virtual machines, so the selected option was **Change datastore**.

Destination disk types in Storage vMotion

During a migration with Storage vMotion, virtual disks can be transformed from thick-provisioned to thin-provisioned or from thin-provisioned to thick-provisioned. The following format options are available:

- Same as Source

Use the format of the original virtual disk. If this option is selected for a raw device mapping (RDM) disk in either the physical or virtual compatibility mode, only the mapping file is migrated.

- Thin-provisioned

Use the thin format to save storage space. The thin virtual disk only uses as much storage space as it needs for its initial operations. When the virtual disk requires more space, it can grow in size up to its maximum allocated capacity. This option is not available for RDMs in physical compatibility mode. If this option is selected for a virtual compatibility mode RDM, the RDM is converted to a virtual disk. RDMs converted to virtual disks cannot be converted back to RDMs.

- Thick-provisioned

Allocate a fixed amount of hard disk space to the virtual disk. The virtual disk in the thick format does not change its size and, from the beginning, occupies the entire datastore space provisioned to it. This option is not available for RDMs in physical compatibility mode. If this option is selected for a virtual compatibility mode RDM, the RDM is converted to a virtual disk. RDMs converted to virtual disks cannot be converted back to RDMs.

Disks are converted from thin to thick format or thick to thin format only when they are copied from one datastore to another. If a disk is left in its original location, the disk format is not converted, regardless of the selection made.

There is also a difference in the way Storage vMotion operates with virtual disks and RDMs.

- For virtual disks in persistent mode, the entire virtual disk is migrated.
- For RDMs in virtual mode, it is possible to migrate the mapping file or convert to thick-provisioned or thin-provisioned disks during migration, as long as the destination is not an NFS datastore.
- For RDMs in physical mode, only the mapping file is migrated, the RDM does not move.

Refer to [Storage vMotion requirements and limitations](#) for more information about requirements and limitations.

Using VMware Storage vMotion

The VMware vCenter Migration wizard can be used to migrate a powered-on virtual machine from one host to another, using vMotion technology. To relocate the disks of a powered-on virtual machine, the virtual machine is migrated using Storage vMotion.

Both of these migration methods can be executed from the same **Migrate** option on a virtual machine as follows:

1. In the vSphere client, right-click the virtual machine for all available options, as shown in Figure 3.




Figure 3: Selecting the Migrate option

The **Select Migration Type** screen is displayed, with the option to either move the virtual machine to another host or to move the virtual machine's storage to another datastore, as shown in Figure 4.

When the virtual machine is powered on, **Change datastore** is the default datastore migration type. The basic settings provide the ability to migrate all of the storage to a single datastore only.




Figure 4: Selecting the Migration Type

2. To display a list of available target datastores, that enable the mapping of individual virtual machine disk (VMDK) files to a specific datastore, select the **Advanced** option as shown in Figure 5.




Figure 5: Selecting the Advanced option

When selecting the destination datastore, as shown in Figure 6, it is important to consider the placement of individual virtual disks when dealing with I/O-intensive applications. Certain components of the application, such as database, logs, and indexes, may perform better or be better protected from component failure, if placed on separate storage devices.




Figure 6: Selecting the datastore

When all of the relevant source virtual disks have been mapped to their respective target datastores, the **Ready to Complete** screen provides a final summary of the selected Storage vMotion job as shown in Figure 7.




Figure 7: Reviewing the summary screen

As can be seen from Figure 7, in this scenario, the OS device was set to remain on its NFS storage and the remaining eight virtual disks were set to be moved to the TARGET_SVMOTION_EXCHANGE datastores.

For more information about using the Migration wizard, refer to the *vSphere Datacenter Administration Guide*.

Storage vMotion requirements and limitations

A virtual machine and its host must meet resource and configuration requirements for the virtual machine disks to be migrated with Storage vMotion.

Storage vMotion is subject to the following requirements and limitations:

- Virtual machines with snapshots cannot be migrated using Storage vMotion. To migrate these machines, the snapshots must be deleted or reverted.
- Virtual machine disks must be in persistent mode or must be RDMs. For virtual compatibility mode RDMs, it is possible to migrate the mapping file or convert to thick-provisioned or thin-provisioned disks during migration, as long as the destination is not an NFS datastore. For physical compatibility mode RDMs, it is only possible to migrate the mapping file.
- The migration of virtual machines during VMware Tools installation is not supported.
- The host on which the virtual machine is running must be licensed for either the Enterprise or Enterprise Plus editions to execute a Storage vMotion operation
- ESX/ESXi 3.5 hosts must be licensed and configured for vMotion. ESX/ESXi 4.0 and later hosts do not require vMotion configuration to perform migration with Storage vMotion.
- The host on which the virtual machine is running must have access to both the source and target datastores.
- A particular host can be involved in up to two migrations with vMotion or Storage vMotion at one time.
- VMware vSphere supports a maximum of eight simultaneous vMotion, cloning, deployment, or Storage vMotion accesses to a single VMFS3 datastore, and a maximum of four simultaneous vMotion, cloning, deployment, or Storage vMotion accesses to a single NFS or VMFS2 datastore. A migration with vMotion involves one access to the datastore. A migration with Storage vMotion involves one access to the source datastore and one access to the destination datastore.

Storage hardware acceleration

Through the use of VMware vStorage APIs for Array Integration (VAAI)—Full Copy, it is possible to accelerate Storage vMotion using compliant storage hardware, enabling the host to offload specific virtual machine and storage management operations to the hardware layer. With storage hardware assistance, the host performs these operations faster and consumes less CPU, memory, and storage fabric bandwidth.

Note VAAI licensing requires the Enterprise edition or higher.

The Full Copy feature offloads the cloning operations to the storage array. The host issues the EXTENDED COPY SCSI command to the array and directs the array to

copy the data from the source LUN to a destination LUN, or to the same source LUN, if required, depending on how the VMFS datastores are configured on the relevant LUNs. The array uses its efficient internal mechanism to copy the data and confirms **Done** to the host. Figure 8 shows how the storage hardware acceleration process is managed.




Figure 8: Storage hardware acceleration with VAAI Full Copy

Full Copy (or VAAI) enables arrays to make copies of certain virtualization objects within the array, without the need to have the ESX server read and write those objects.

To benefit from the hardware acceleration functionality, you must have:

- ESX version 4.1 or later
- A storage array that supports hardware acceleration (for example, CLARiiON FLARE 30)

On the VMware vSphere Server, hardware acceleration is enabled by default.

To change this setting, go to **ESX Server Configuration Tab > Software – Advanced Settings > DataMover > DataMover.HardwareAcceleratedMove**, as shown in Figure 9, where:

- 0 = disabled
- 1 = enabled




Figure 9: Hardware acceleration settings

To enable hardware acceleration, run the following command on the ESX version 4.1 console:

```
esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedMove
```

A required configuration step when using a CLARiiON array that supports the Full Copy/Array Accelerated Copy feature: the ESX host initiator records must be configured using **failovermode 4**, that is, asymmetric logical unit access (ALUA) mode on the CLARiiON.

The Full Copy feature is only supported when the source and destination LUNs belong to the same storage array. Currently, it is not supported for cross-array migrations.

For more information on VAAI, visit the VMware Knowledge Base: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1021976

Comparing Storage vMotion with and without VAAI enabled

In the test environment, a single Windows virtual machine running Microsoft Exchange 2010 was used. The virtual machine’s boot device was on NFS storage, so the goal was to migrate the application data only. Apart from the boot device, the virtual machine had eight virtual disks, spread evenly across two 100 GB datastores that were configured on two separate FC LUNs as shown in Table 4.

Table 4: Datastore configuration

Datastore 1 (100 GB)	Datastore 2 (100 GB)
DB-01-Database (25 GB)	DB-01-Logs (10 GB)
DB-02-Logs (10 GB)	DB-02-Database (25 GB)
DB-03-Database (25 GB)	DB-03-Logs (10 GB)
DB-04-Logs (10 GB)	DB-04-Database (25 GB)

The Storage vMotion operation, although executed against a single virtual machine, required the simultaneous migration of eight separate virtual disks from two source datastores to two separate target datastores.

Figure 10 displays the disk activity, as seen through the ESX storage adapter counters (**vmhba0** and **vmhba1**), within the vSphere Client performance views for both VAAI and non-VAAI Storage vMotion operations.




Figure 10: Disk activity

These Storage vMotion operations were conducted with the virtual machine powered up but not under any load. As can be seen when VAAI is enabled, the average commands per second are vastly reduced, from a combined total of approximately 3,600 commands per second down to approximately 800 commands per second. This drop in host bus adapter (HBA) activity is directly attributable to the fact that, with VAAI enabled, ESX no longer needs to conduct as many inter-datastore operations because the array manages this data transfer on the back end.

While the VAAI-enabled migration was slightly quicker, it is noticeable that the profile of the read activity was completely different when VAAI was enabled. There was a short spike in the read commands per second at the beginning of the VAAI-enabled Storage vMotion, which then ceased for the remainder of the migration.

This is the nature of hardware offloading, as the ESX server reads the full contents of the source device once and sends it to the array. The array migrates the data to the target device and, on completion, sends the **Done** command to the ESX server.

Impact on the CLARiiON storage processors

It is also interesting to note the impact this offloading has on the CLARiiON throughput, as shown in Figure 11. Both datastores were configured on separate storage processors (SPA and SPB) so that the load is spread evenly across both.




Figure 11: CLARiiON throughput comparison

The array statistics correlate perfectly with what the ESX server observes during the online Storage vMotion operations. With the traditional, non-VAAI Storage vMotion operation, the array detects far more read and write activity from the ESX server. With VAAI enabled, this activity greatly decreases because, once the data is copied to the array, the hardware offloading takes care of the subsequent copy operations to the source device at the back end.

It is interesting to note that the CLARiiON storage processors are actually busier when VAAI is not enabled, as shown in Figure 12. So instead of the hardware offloading putting an increased load on the CLARiiON storage array, it did the opposite in this case; the array was able to use its own internal efficiencies for the back-end copy rather than servicing front-end I/O with non-VAAI Storage vMotion operations.




Figure 12: CLARiiON SP utilization comparison

Impact on Exchange 2010

Neither the virtual machine nor the ESX server was short of CPU or memory resources during the Storage vMotion operations, but the associated increase in I/O activity could have had an adverse effect on the response times of the application.

To validate Exchange 2010 performance, Microsoft LoadGen was run for two 2-hour periods, once with VAAI disabled and once with VAAI enabled, with the Storage vMotion operation executed 30 minutes into each test.

While the final results in terms of the overall IOPS achieved and the average response times for databases and logs were the same for each run, there was a noticeable increase in the database seconds per read (DB sec/Read) response times for the duration of the Storage vMotion operations as shown in Figure 13.




Figure 13: Exchange 2010 latencies comparison

The baseline figures reflect what the normal latencies were at either side of the Storage vMotion operations. In terms of the impact to Exchange, there was no distinguishable difference in read and write latencies during VAAI and non-VAAI migrations. Both types of migration added an extra three microseconds to the DB sec/Read, and one extra microsecond to the database seconds per write (DB sec/Write). The log writes were unaffected by the migrations. This is not to say that Storage vMotion has no effect on the log devices, simply that the impact was not measurable in this case, since the sequential nature of Exchange log devices lends itself to good performance when the underlying storage is correctly configured.

Conclusion

VMware Storage vMotion is possibly the easiest and most convenient method of migrating storage in a virtualized environment. It is one of the few methods that provides the VMware administrator with online, nondisruptive mobility across the underlying storage infrastructure. Once the vSphere Server detects both the source and target datastores, it is simply a matter of scheduling the migration process itself. Storage vMotion is now further enhanced by VAAI with Hardware Offloading and Full Copy functions that dramatically reduce the server workload required to complete the migration using storage array technology and efficiencies.

Data migration with VMware vCenter Converter

VMware vCenter Converter

Two versions of the VMware vCenter Converter are available: standalone and integrated. The VMware vCenter Converter Standalone used in this test scenario was version 4.0.1 (build 161434). The VMware vCenter Converter module tested was the version integrated into VMware vCenter 4.1.

Both tools allow for hot cloning (converting the powered-on machine) or cold cloning (booting from the VMware vCenter Converter boot CD), regardless of whether the source is a physical or virtual machine.

While in most scenarios VMware vCenter Converter is used to convert a physical host or a virtual machine from a different hypervisor to a VMware virtual machine, it is often considered as a potential candidate tool for the migration of a VMware virtual machine from one place to another.

One distinction should be considered in terms of using VMware vCenter Converter for migrations. Strictly speaking, VMware vCenter Converter copies and clones the virtual machine to another location, rather than migrating or moving it. A second, separate machine is created that retains all of the operating system and data characteristics.

When converting from physical machines or non-VMware hypervisor machines, it is assumed that the resulting VMware virtual machine will have a different set of hardware. This is partially the case also if the VMware vCenter Converter is used to migrate an existing VMware virtual machine to another VMware virtual machine—for example, resulting in different MAC addresses, which could potentially have consequences for any MAC address-based licensing or network security that is already in place.

The following scenarios were tested for the standalone and integrated versions of VMware vCenter Converter:

- Hot cloning with synchronization and switch features
- Cold cloning

Hot cloning with synchronization and switch features

Hot cloning is supported by both the VMware vCenter Converter Standalone and the VMware vCenter Converter module (vCenter plug-in). This enables a copy of the running machine (physical or virtual) to be created on a VMware virtual machine running on a vSphere host, without interruption of service. Note that there is a distinction between creating a copy without interruption, and transferring production without interruption. VMware vCenter Converter Standalone supports the former, but not the latter. The usefulness of VMware vCenter Converter as a migration tool is therefore limited if its synchronization and switch features are disabled.

- With the synchronization feature on, VMware vCenter Converter performs an initial copy of all requested volumes, and then does an incremental resynchronization of the data that was changed during the initial copy window. In an active system, data is continually changing. The resynchronization feature executes only once, so if data is still being changed during resynchronization, then this data is not captured.

To provide for this, VMware vCenter Converter enables the user to select services (in the case of Windows) that should be shut down prior to performing the resynchronization, thereby preventing data loss.

- With the switch feature on, VMware vCenter Converter powers down the source machine, and powers up the newly created virtual machine, in addition to performing any requested customization.

Configuring a hot cloning operation

The VMware vCenter Converter Standalone interface is similar to the VMware vCenter Converter integrated interface, which is used to illustrate the following hot cloning configuration process.

Before starting the configuration:

- Install the VMware vCenter Converter module
- Install the plug-in installed on the VI Client

Use the following steps to configure the hot cloning operation, as done in this scenario:

1. Select **Import Machine** from the context menu of the import destination (in this case, a vSphere host) as shown in Figure 14.




Figure 14: Selecting Import Machine

The **Source System** screen is displayed as shown in Figure 15.




Figure 15: Source System screen

2. Select **Powered-on Machine** as the source type and provide the relevant credentials for the powered-on machine: **IP address or name**, **User name**, **Password**, and **OS Family**.
3. Click **Next** to continue.

A message is displayed as shown in Figure 16.




Figure 16: Uninstall options

To complete the import, a VMware vCenter Converter agent is temporarily installed on the source machine. These files can be automatically or manually removed after the import, depending on the method selected.

4. Select the automatic uninstall method and click **Yes** to continue.

The **Destination Location** screen is displayed as shown in Figure 17.




Figure 17: Destination Location screen

5. Select a datacenter from the **Inventory** to hold the destination virtual machine and provide the relevant names for the **Virtual machine name**, **Datastore**, and **Virtual machine version**.

Note If the source is an existing virtual machine in the same vCenter instance, then specify a new name for the virtual machine.

6. Click **Next**.

The **Options** screen is displayed as shown in Figure 18 where parameters can be configured for the conversion.




Figure 18: Options screen

7. In the **Data to Copy** section, from the **Advanced** section, select target datastores individually for each volume in the source machine.

Note It is also possible to convert to thin format at this point, if required.

8. Ensure that the correct virtual network on which to run the machine is selected as shown in Figure 19.

VMware vCenter Converter defaults to the first alphabetical network name (and not necessarily the previous network used if the source is a VMware virtual machine).




Figure 19: Options screen - selecting the correct virtual network

9. If resynchronizing the data after the initial copy, select the services you want to stop on the source machine before starting to resynchronize. This is to ensure that no data is lost before switching over to the new target machine.

In this scenario, the SQL server is running actively on the source and the SQL services are stopped before performing the final resynchronization as shown in Figure 20.




Figure 20: Options screen - stopping SQL services before resynchronization

Note This scenario also uses advanced options, including those to resynchronize and power off the source as well as to power on the destination machines.

Since an existing virtual machine is being migrated in this scenario, there is no need to install VMware tools as these were already present.

10. Click **Next**.

The **Summary** screen is displayed showing all of the selected options prior to cloning as shown in Figure 21.




Figure 21: Summary screen

Performance impact of hot cloning

During a hot cloning operation, the source disks are copied to another location. This results in a substantial increase in read I/O on the source LUNs. This may or may not have an impact on the source system performance, depending on the underlying configuration of the source LUNs.

In this case, there was no measurable impact to the performance of the SQL Server response times or the transactions per minute executed by the SQL DVD Store application. However, the additional read I/O was very apparent when looking at the utilization of the source LUNs, as demonstrated in the graph in Figure 22.




Figure 22: LUN utilization

Cold cloning

Cold cloning creates a copy of the virtual machine on the target VMware vSphere system while the source machine (physical or virtual) is shut down. More accurately, this means a virtual machine that is powered off, or a physical machine that is not running its installed operating system but has instead been booted from a VMware vCenter Converter Boot CD.

The length of time taken to complete the copy (and therefore the level of disruption to service) is therefore entirely dependent on the quantity of data to be transferred, as well as the specification of the source and target hardware. The advantage of a cold cloning process is that there is no chance of data being updated on the source system, so there is no need for incremental resynchronization as in the hot cloning process.

There is no performance impact to measure since the source system is inactive during the cloning process.

As with the hot cloning process, the target system is a copy of the source system, but with a different set of similar hardware, so there are different MAC addresses and so on, on the target virtual machine.

The process of cold cloning is very similar to hot cloning, except for the step where the type of source machine is selected. Previously, a powered-on virtual machine was selected but, in this scenario, one of the other options was selected as appropriate, depending on whether the source is an existing vSphere virtual machine, a virtual machine from an alternate hypervisor, or a physical machine.

Data migration with EMC VPLEX Metro

Introduction to the EMC VPLEX family

The EMC VPLEX family, with the EMC GeoSynchrony™ operating system, makes it possible for users to overcome the physical barriers of data centers, enabling them to access data for read and write operations at different geographical locations concurrently. This is achieved by synchronously replicating data between data centers, while depending on the hosts accessing the storage devices to manage consistency through the use of intelligent, distributed lock management.

This capability, in a VMware context, enables functionality that was not previously available. Specifically, the ability to concurrently access the same set of devices, independent of the physical location, enables geographical vMotion, based on the VMware virtualization platform. This allows for transparent load sharing between multiple sites, while providing the flexibility of migrating workloads between sites in anticipation of planned events, such as hardware maintenance.

Furthermore, in case of an unplanned event that causes disruption of services at one of the data centers, the failed services can be quickly and easily restarted at the surviving site with minimal effort. The capabilities in VPLEX are a strong complement to existing DR solutions such as VMware SRM.

The VPLEX family consists of two products: VPLEX Local and VPLEX Metro.

- VPLEX Local provides simplified management and nondisruptive data mobility across heterogeneous arrays within a data center. With a unique scale-up and scale-out architecture, the VPLEX system's advanced data caching and distributed cache coherency provide workload resiliency, automatic sharing, balancing, and failover of storage domains, and enables local data access with predictable service levels.
- VPLEX Metro delivers distributed federation capabilities and extends access between two locations at synchronous distances. VPLEX Metro leverages AccessAnywhere™ storage that allows data to be moved, accessed, and mirrored transparently between data centers, effectively allowing storage and applications to work between data centers as though those physical boundaries were not there.

Figure 23 contrasts the difference in architecture between the traditional method of SAN-based storage access and VPLEX storage, which presents storage through a virtualization layer.




Figure 23: Architectural differences between SAN-based and VPLEX storage

Data migration with VPLEX

Using the data mobility functionality of VPLEX Local and VPLEX Metro, IT organizations can seamlessly migrate data between storage tiers or from lease arrays, with no disruption in service.

Also, in a virtual environment, IT organizations can use the VPLEX Metro configuration during data center maintenance operations or in a disaster avoidance scenario. With access to a virtual machine’s configuration files from both clusters residing in different data centers, a virtual machine can be migrated across data centers. In such a scenario, where virtual machines can be migrated between sites, VPLEX provides a robust business continuity solution.

Note VPLEX is a disaster avoidance solution and not a disaster recovery (DR) solution. With VPLEX, virtual machines can be migrated live and online with no user impact, whereas in a DR solution there is an impact on availability during the failover process.

Configuration of VPLEX Metro

Figure 24 illustrates the configuration of VPLEX Metro that enables live migration of virtual machines between two sites, separated by distance.




Figure 24: VPLEX Metro enabling VMware vMotion of applications across sites

As shown in Figure 24, each site has a VPLEX cluster with access to physical storage. The cluster at each site communicates with the other clusters through FC. The FC extension between the VPLEX clusters can be either with dark fiber extending between the VPLEX clusters or, from GeoSynchrony version 4.0 and upward, with an FC over IP (FCIP) tunnel on the IP WAN between the data centers.

Figure 24 also shows how the federation capability of VPLEX Metro enables the creation of a distributed volume that has the same SCSI identification, independent of the location from which the device is accessed. Therefore, the two VMware ESX hosts consider the distributed volume as the same device and enable capabilities, such as VMware vMotion, that were traditionally available only in a single data center.

Prior to VPLEX, migrating a virtual machine from site to site required a stretched SAN. The downside was that both VMware VMotion and VMware Storage vMotion were required to make the storage available at the distant site. This significantly adds to the time needed for failover, which limits the ability to easily move virtual machines from site to site, as can be done using VPLEX.

VPLEX includes the following features:

- VPLEX provides the ability to share SAN storage across distances up to 100 km (or less than 5 microseconds (ms) of round trip latency).
- VPLEX supports clustered file systems, such as those from VMware and Microsoft, which enhance the ability to utilize the functionality across distance.

- VPLEX supports block-level storage virtualization over FC with no support for iSCSI or NFS.
- The disparate distance between sites is seamless, so neither the servers nor the applications have knowledge of VPLEX in the storage environment.

VPLEX virtual volumes

The virtual volume that VPLEX presents to a host contains both extents and devices, as shown in Figure 25. The ability to migrate these extents and devices while keeping the unique identifier network address authority (NAA) of the virtual volume is what makes the migration seamless and transparent to the end host.




Figure 25: Virtual volume

Introducing VPLEX into an existing environment

To deploy a new VPLEX into an already existing VMware environment, several approaches can be taken. The simplest, most nondisruptive method is to present new storage to the VPLEX and then present the virtual volume to an ESX host. Deployment can be completed by using VMware Storage vMotion to move the virtual machines onto the new VPLEX datastore. In this scenario, there is no interruption in service, making this method completely nondisruptive.

However, if network bandwidth is a constraint and the goal is to migrate the content of an entire datastore, complete with virtual machines and ISO images, then the migration is disruptive to the virtual machines residing on the datastore in question.

The required steps, using this disruptive method, are:

1. Shut down the virtual machines on the datastore to be migrated to the VPLEX.
2. Perform the necessary LUN masking on the storage array.
3. Claim the LUN by the VPLEX and present it to the ESX host as a virtual volume.
4. Rescan the ESX HBAs.
5. Add storage while keeping the existing VMFS signature.
6. Add the virtual machines to the inventory.

Migration scenario

In the following migration scenario SQL and Exchange 2010 hosts are migrated nondisruptively between data centers, using VPLEX and VMware vMotion.

VPLEX Metro configuration

From the VPLEX Management Console, the SQL and Exchange 2010 database and log LUNs are configured as distributed RAID 1 (DR1) devices. These devices, whose mirrors are in two geographically dispersed locations, are shown in Figure 26.




Figure 26: VPLEX Management Console

The distributed devices are composed of storage located on both clusters in the VPLEX Metro configuration. In this scenario, each cluster has a CLARiiON CX4-480 presenting storage to the VPLEX that makes up the distributed devices. Each virtual volume presented to both clusters has the same unique identifier NAA with the same VMFS signature. Therefore, the virtual machines and their VMDK files can be accessed on either cluster, enabling vMotion operations as opposed to Storage vMotion. This is illustrated in Figure 27 and Figure 28.




Figure 27: VPLEX-Cluster-A devices




Figure 28: VPLEX-Cluster-B devices

Table 5 provides a summary of the virtual disks used by the virtual machines in this scenario.

Table 5: Virtual disks in use

Capacity	Number of LUNs	RAID type
11 GB	3	RAID 5 (4+1)
101 GB	2	RAID 5 (4+1)
1 TB	1	RAID 5 (4+1)

VMware cluster configuration

The two VMware clusters (VPLEX-Cluster-A and VPLEX-Cluster-B) each contain a single ESX 4.1 host. The Exchange-2010 vApp, which consists of an Exchange 2010 server virtual machine and a Windows 7 virtual machine running Microsoft Exchange Load Generator (LoadGen), is configured in Cluster A, while the SQL-DVD-Store virtual machine is configured in Cluster B. The VPLEX-Cluster-A configuration is shown in Figure 29.

Each ESX server has six LUNs (VPLEX distributed devices) allowing both clusters to have access to the same NAA IDs.




Figure 29: VPLEX-Cluster-A configuration

Exchange 2010 performance with vMotion over distance

To measure any potential performance impact on the Exchange 2010 application during the vMotion over distance operation, a baseline performance test was required.

This test consisted of a 2-hour LoadGen run, as shown in Figure 30. This baseline test was run locally on Cluster A in Datacenter A, with no other operations in progress.




Figure 30: Baseline test

The baseline test results were then compared to the test results observed while the vMotion over distance operation was in progress between Cluster A and Cluster B.

The end result was that there were no significant performance differences in Exchange 2010 counters between the baseline test and the vMotion over distance test. Importantly, both the Exchange 2010 virtual machine and the Windows 7 virtual machine were migrated to Cluster B in Datacenter B, with no interruption of operations.

The results of the 2-hour LoadGen test are detailed in Figure 31.

\\Exchange-2010	Baseline	Storage vMotion
LogicalDisk	_Total	_Total
Avg. Disk sec/Read	0.004	0.003
Avg. Disk sec/Write	0.002	0.002
Disk Reads/sec	56.926	50.615
Disk Transfers/sec	388.710	404.234
Disk Writes/sec	331.784	353.620
MSExchangeIS		
RPC Averaged Latency	1.273	2.969
RPC Operations/sec	1,899,460	1,949,080
Processor	_Total	_Total
% Processor Time	51.808	53.327

Figure 31: 2-hour LoadGen test results

Note **MSExchangeIS RPC Averaged Latency** increased by 100 percent but this is well within the 50 ms threshold for greater than 3,000 users. For more information on remote procedure call (RPC) performance counters, visit the Microsoft Technet website: [http://technet.microsoft.com/en-us/library/aa998266\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/aa998266(EXCHG.80).aspx)

The Exchange 2010 and Windows 7 virtual machines were migrated from Cluster A (Datacenter A) to Cluster B (Datacenter B) using vMotion over distance. Both the Exchange 2010 virtual machine and Windows 7 virtual machine remained online for the duration of migration.

For more information, refer to the white paper *Using VMware Virtualization Platforms with EMC VPLEX—Best Practices Planning*.

Conclusion

The use of EMC VPLEX storage ensures that the performance of the application is almost the same at both data centers. In addition, an Active/Active data center proves to be an operational reality. Finally, migration to a remote data center is feasible not only from a technical perspective (application mobility is possible) but also from a business standpoint (application performance is not adversely affected).

Data migration with CLARiiON SAN Copy

Introduction to EMC CLARiiON SAN Copy

EMC CLARiiON SAN Copy is a storage-system-based application that is available as an optional package. SAN Copy is designed as a multipurpose replication product for data mobility, migrations, content distribution, and disaster recovery.

SAN Copy enables the storage system to copy data at a block level directly across the SAN, from one storage system to another, or within a single CLARiiON system. While the software runs on the CLARiiON storage system, it can copy data from, and send data to, other supported storage systems on the SAN.

It is important to note that the migration of production data with SAN Copy is a disruptive operation. Whereas some of the other technologies discussed in this white paper provide nondisruptive migrations, SAN Copy is disruptive as the systems must manually switch between accessing the old and the new copy. SAN Copy supports operations over both FC and iSCSI and can be managed through Unisphere, the Navisphere® command line interface (NaviSecCLI), or Replication Manager. Be aware, however, that SAN Copy does not provide the complete end-to-end protection that products such as MirrorView™, RecoverPoint, and Symmetrix Remote Data Facility (SRDF®) provide.

Environment topology

Figure 32 illustrates the environment topology for the SAN Copy data migration scenario.




Figure 32: SAN Copy data migration

Uses for SAN Copy

The common uses for SAN Copy are:

- **Data mobility**—eliminate impact on production activities during data mobility tasks
- **Data migration**—easily migrate data from qualified storage systems to the CLARiiON system
- **Content distribution**—regularly push updated production data to remote locations
- **Disaster recovery**—protect and manage applications through integration with Replication Manager

For the purpose of this white paper, the focus is on using SAN Copy for performing data migrations in a VMware environment.

Migrations available with SAN Copy

With SAN Copy, the following data migrations are possible:

- Migration between CLARiiON arrays
 - Migration within the same CLARiiON array
 - Migration between CLARiiON, Symmetrix, and third-party arrays
-

Benefits of SAN Copy

The main benefits of SAN Copy are:

- Optimal performance, as data is copied directly through the SAN
 - No host resources are required, as SAN Copy executes on the storage array
 - Interoperability with many heterogeneous storage systems
-

SAN Copy full and incremental sessions

SAN Copy supports two types of sessions:

- Full SAN Copy

A full session copies the entire contents of the source LUN to the destination LUN(s) every time the session is executed. Full sessions can be a push or a pull with any qualified storage system.

- Incremental SAN Copy

An incremental session requires a full copy only once, which is referred to as an initial synchronization. Each session after that copies only the changed data from the source LUN to the destination LUN. For incremental sessions, the source LUN must reside on a CLARiiON system but the destination LUN can reside on any qualified storage system.

SAN Copy in a VMware environment

Since it is array-based, the focus of SAN Copy is entirely on the source LUN and its contents. From an operational perspective, the VMware layer is irrelevant to SAN Copy. Regardless of whether a LUN is formatted as a VMFS datastore containing multiple virtual disks, or a LUN is passed through to a virtual machine as an RDM, SAN Copy copies every block of data within that source LUN to the destination LUN.

If using Full SAN Copy to migrate a SAN LUN containing a VMFS datastore, the VMFS datastore must be quiesced before the migration can take place to ensure a consistent point-in-time image of the data. As this migration would be a disruptive once-off operation, the simplest method to quiesce the VMFS datastore is to shut down the virtual machines currently residing on and accessing the datastore. Of course, it is also possible to use Replication Manager to quiesce the VMFS datastore. Replication Manager is normally used in the event of regular (daily and weekly), incremental SAN Copy sessions where the production systems need to remain online while SAN Copy sends updated data to the destination LUN.

Requirements and considerations

SAN Copy is bound by the following requirements and considerations.

- **Data consistency**—there is a difference in the requirements for quiescing the source LUN when using full and incremental SAN Copy sessions. The source LUN must be quiesced for the duration of a full SAN Copy session. For an incremental SAN Copy session, the source LUN only needs to be quiesced just before it begins. Quiescing the source LUN can be achieved by using the Navisphere **admhost** utility on Windows machines to flush the file systems, or in the case of a VMFS datastore, shutting down the virtual machines residing on the datastore achieves the same goal. If quiescing the source LUN is not possible, then the destination image is crash-consistent on completion of the SAN Copy session.
- **Application consistency**—Replication Manager can be used with SAN Copy to ensure application consistency. Replication Manager supports Microsoft's virtual shadow copy service (VSS) architecture, which allows for the creation of hot, point-in-time Exchange images, without disrupting the production server. SQL, SharePoint, Oracle, and DB2 are also supported (including their relevant hot and online backup modes) as well as VMware VMFS datastores and Hyper-V guests that use iSCSI and CLARiiON or Celerra storage.
- **Source LUNs and incremental SAN Copy support**—CLARiiON storage systems leverage EMC SnapView™ functionality to execute incremental SAN Copy sessions. For this reason, the source LUN must reside on a CLARiiON storage system.

Migration duration and performance

SAN Copy enables users to throttle the speed and performance of the SAN Copy session as shown in Figure 33. The throttle values range from 1 to 10 (1 is slowest, 10 is fastest, and the default is 6).




Figure 33: Setting the throttle value

With the throttle value set to 10, the system attempts to get the data from the source to the target as quickly as possible, utilizing all available resources. As a result, storage system resources are used far more aggressively than would be the case if the throttle value was set to 1. With throttle value set to 1, the data migration takes the maximum amount of time to complete but uses the lowest amount of storage system resources. The throttle value can be changed dynamically during an active SAN Copy session.

SAN Copy also enables the user to optimize performance by providing the ability to control latency and link bandwidth as shown in Figure 34.




Figure 34: Setting the link bandwidth value

These values are used by SAN Copy to calculate buffer space for optimal performance.

As with any migration, due consideration must be given to the potential impact on other applications and shared resources during the migration process.

CLARiiON SAN Copy setup overview

The high-level steps required for the successful setup and execution of a SAN Copy session are to:

1. Identify the source and target storage arrays and LUNs.
2. Configure SAN connectivity.
3. Select the source LUN.
4. Select the target LUN.
5. Verify SAN connectivity.
6. Configure the SAN Copy session properties.
7. Save the SAN Copy session.
8. Execute the SAN Copy session.

The most important step is to configure the end-to-end SAN connectivity between the source and the target LUNs. This part of the process is crucial. When migrating from one CLARiiON array to another, zone the CLARiiON storage SP ports of the source array to the SP ports of the target array, and enable the source array to access the relevant LUNs on the target array.

Figure 35 illustrates the LUN masking procedure on the target CLARiiON (**Springfield-CX480**), allowing access to the source CLARiiON (**Quahog-CX480**), as managed by Unisphere.




Figure 35: LUN masking procedure

Table 6 details the steps needed to complete the procedure.

Table 6: Completing the procedure

Step	Action
1	Right-click on the CLARiiON storage group containing the LUN.
2	Select the port to be granted access.
3	Click OK . The port is then connected to that storage group as shown in Figure 35.

Migrating to third-party storage arrays

If a storage array other than CLARiiON is being used, the zoning requirements are the same but the LUN masking must be configured both inside and outside the control of Unisphere. The SAN Copy session is still controlled by Unisphere. The world wide name (WWN) of the target LUN and the port WWN must be entered in the CLARiiON hosting the SAN Copy session, so that the CLARiiON can recognize the identity of the LUN in question. This is illustrated in Figure 36.




Figure 36: Entering the WWN

The third-party storage array must then allow LUN access to the WWN of the CLARiiON SP port. This end-to-end connectivity is central to the entire process and is verified during the setup of the SAN Copy session.

The following arrays are qualified to host SAN Copy operations:

- Older CLARiiON CX arrays such as CX700, CX600, CX500, and CX400
- All CLARiiON CX4 arrays
- All CLARiiON CX3 arrays

SAN Copy-compatible storage systems (that is, systems that can participate in but not host SAN Copy sessions) include:

- Symmetrix 8000 series, DMX-2, DMX-3, and VMAX
- Third-party vendors including:
 - HP—EMA, EVA, and XP arrays
 - IBM—DS and Fast arrays
 - HDS—99xx, TagmaStore, and Thunder arrays
 - 3PAR—E200, S400, and S800 arrays
 - Sun—StorEdge 9990 and T3/T3+

Note The SAN Copy-compatible storage systems listed are correct at the time of publication. Refer to the *EMC Support Matrix* for the full list of supported storage arrays.

Migration scenario

In this scenario, Microsoft Exchange 2010 is migrated from site to site.

As SAN Copy is a disruptive procedure, the process of migrating the Exchange 2010 server and storage from site to site requires application downtime. The options, in this case, are to minimize the amount of SAN Copy operations by using full SAN Copy sessions or to minimize the downtime through incremental SAN Copy. For the incremental option, the bulk of the data is copied prior to the downtime. Then once the environment is quiesced and is ready to migrate, the incremental changes can be copied to the target. This scenario uses incremental SAN Copy sessions.

The environment to be migrated is a single Exchange 2010 Mailbox server virtual machine with two 100 GB datastores, each containing four VMDK files for Exchange data and logs, as shown in Table 7.

Table 7: Datastores

Datastore 1 (100 GB)	Datastore 2 (100 GB)
DB-01-Database (25 GB)	DB-01-Logs (10 GB)
DB-02-Logs (10 GB)	DB-02-Database (25 GB)
DB-03-Database (25 GB)	DB-03-Logs (10 GB)
DB-04-Logs (10 GB)	DB-04-Database (25 GB)

Both the source storage and the target storage arrays are CLARiiON CX4-480 storage arrays, so the entire operation is managed within Unisphere. The source CX4-480 is the SAN Copy session owner that pushes the data to the target CX4-480 as shown in Figure 37.

Migrating to a separate VMware vSphere cluster is also included in this test. It is possible to present the target LUNs back to the original cluster, if preferred, once the data is copied. Be aware of the potential for datastore resignaturing in that situation.




Figure 37: Migrating from source to target

The setup and preparation for the SAN Copy operations were independent of vSphere. As previously stated, the end-to-end connectivity is central to SAN Copy operations. Once the relevant SP ports from **Quahog-CX4-480** were zoned to **Springfield-CX4-480**, the appropriate LUN masking could take place internally.

The LUNs to be migrated had CLARiiON LUN IDs of LUN 22 and LUN 23. These LUNs were renamed to SANCopy_Exchange_DS1_LUN22 and SANCopy_Exchange_DS2_LUN23 within Unisphere.

Creating a SAN Copy session

To begin the migration setup, select the source LUNs by launching the **SAN Copy Wizard** in Unisphere. Navigate to the **Replicas** section where the **SAN Copy Wizard** can be started from the options pane at the left side of the screen, as shown in Figure 38.




Figure 38: Launching the SAN Copy Wizard in Unisphere

The **SAN Copy Wizard** guides you through the following stages:

- Selecting the storage array that owns the session
- Defining the type of session as a full or incremental SAN Copy
- Selecting the source and destination devices
- Saving the session for future use and execution

After creating and saving the SAN Copy session for each of the Exchange datastores, it is possible to edit or change the session properties at any stage before the session begins, as shown in Figure 39.




Figure 39: Changing SAN Copy session properties

Session variables, such as the SAN Copy session throttle value and the link bandwidth values, can be changed, if required. It is also possible to edit the SAN connectivity type—FC or iSCSI—as long as the connectivity between the source and target is correct.

Even though this is an incremental SAN Copy session, a full synchronization is initially required. Start by right-clicking on the saved SAN Copy session to display the associated command options, as shown in Figure 40.




Figure 40: Starting full synchronization

Managing and monitoring the SAN Copy session

The SAN Copy session can be managed and monitored in Unisphere, as shown in Figure 41. If required, users have the option to pause and resume the SAN Copy operation, as well as transfer the ownership of the session to the alternate SP.



Figure 41: Monitoring the SAN copy session

Alternatively, SAN Copy sessions can be initiated and monitored using the command line, as shown in Figure 42.

```
C:\Users\Administrator>naviseccli -h 172.30.220.161 sancopy -info -name SANCOPY_EXCHANGE_DS1
Copy Descriptor Name: SANCOPY_EXCHANGE_DS1
Copy Descriptor ID: 4
Owner: SPA
Source LUN WWN: 60:06:01:60:09:60:26:00:10:10:09:BD:69:A1:DF:11
Number of Blocks to Copy: 209715200
Number Of Destinations: 1
Destination LUN WWN: 60:06:01:60:30:50:26:00:A3:ED:F5:34:C3:AA:DF:11
Session Status: Session is active.
Initial Throttle: 10
Current Throttle: 10
Transfer Count: 24347648
Percent Complete: 11
Completion Time: N/A
Failure Status: No Failure

C:\Users\Administrator>naviseccli -h 172.30.220.161 sancopy -info -name SANCOPY_EXCHANGE_DS2
Copy Descriptor Name: SANCOPY_EXCHANGE_DS2
Copy Descriptor ID: 1
Owner: SPB
Source LUN WWN: 60:06:01:60:09:60:26:00:DE:DA:E0:C9:69:A1:DF:11
Number of Blocks to Copy: 209715200
Number Of Destinations: 1
Destination LUN WWN: 60:06:01:60:30:50:26:00:E9:61:22:4E:C3:AA:DF:11
Session Status: Session is active.
Initial Throttle: 10
Current Throttle: 10
Transfer Count: 25371648
Percent Complete: 12
Completion Time: N/A
Failure Status: No Failure
```

Figure 42: Using the command line for SAN Copy sessions

To start the SAN Copy session, use the following commands:

```
naviseccli -h <IP address> sancopy -start -name <session name>
naviseccli -h <IP address> sancopy -start -name <session name>
```

To monitor the SAN Copy session, use the following commands:

```
naviseccli -h <IP address> sancopy -info -name <session name>
naviseccli -h <IP address> sancopy -info -name <session name>
```

The initial copy of the data can be executed without disrupting production operations. The subsequent incremental SAN Copy session is executed after the source or production virtual machine is shut down and removed from the inventory.

The second copy of the data transfers only those blocks of data that changed since the initial copy was created. Therefore, the transfer time is shorter, which improves the recovery time objective (RTO) and enables minimum downtime for the Exchange 2010 server.

From Unisphere, it is possible to view how much data has changed since the full copy was created by navigating to the **Mark** tab in the properties of the SAN Copy session as shown in Figure 43. The *mark* is the point in time at which the image of the LUN was taken.




Figure 43: Viewing the session status from the Mark tab

When the final incremental transfer of the data is completed, all the production data is copied in a consistent state to the destination LUN. For example, the value of the **Number of Blocks Changed to be Copied** shown in Figure 43 changed to "0" after the transfer, as shown in Figure 44.




Figure 44: Viewing data changes after the final transfer

Figure 45 shows how Unisphere displays the completed SAN Copy session.





Figure 45: Viewing a completed session

The session properties also display the bandwidth and latency observed during the session.

When the SAN Copy operations are complete, the destination site must be configured to access the data. The destination site (in this scenario, *Springfield*) has a separate VMware ESX server from the ESX server at the source site, so once the ESX server is zoned and masked to the new storage array, a rescan of the HBAs is required.

This rescan detects the newly assigned datastores. As these datastores already have signatures from the original ESX server, they are resigned by the destination ESX server when they are imported into the inventory, with the prefix “snap-xxx”. These datastores can be renamed with more appropriate labels, if preferred, as shown in Figure 46.



Datastore	Status
NS-40_LUN_30	Normal
NS-40_LUN_40	Normal
snap-25af2278-SANCOPY_EXCHANGE_DS2...	Normal
snap-3403c64e-SANCOPY_EXCHANGE_DS1...	Normal
tce-r900-07	Normal

Figure 46: Newly assigned datastores

To complete the migration process, browse to the relevant datastore, locate the VMX file for the Exchange 2010 virtual machine, and select **Add to Inventory**. Ensure that the relevant network labels and server resources are available on the new site before booting the newly migrated virtual machine.

Note SAN Copy only copies data from the source to the target site. The data still exists on the source site (in this scenario, **Quahog**), both on the ESX server and on the storage array. The appropriate housekeeping must be completed on the source site.

Conclusion

EMC CLARiiON SAN Copy is a valid solution for customers looking to migrate their virtualized environments data from site to site, with the minimum of disruption and impact to their production operations. SAN Copy operates independently of the virtual infrastructure and, as such, removes the requirement to commit any ESX or ESXi resources that would otherwise be necessary to drive a data migration.

For migration scenarios where solutions such as VMware Storage vMotion or proprietary array replication such as MirrorView or SRDF are not suitable, SAN Copy provides a heterogeneous solution that can manage and execute all data movement across the SAN. It also provides an option to migrate from one array to another, without moving virtual machines from one VMware vCenter Server instance to another, as is required by VMware vCenter SRM.

Data migration with CLARiiON LUN Migrator

CLARiiON virtual LUN Migrator

The virtual LUN migration feature on the CLARiiON array provides a simple way of migrating data. CLARiiON virtual LUN Migrator transparently moves data from a source LUN to a destination LUN of the same or larger size within a single storage system. LUN migration can enhance performance or increase storage capacity by enabling users to migrate to a LUN with different characteristics—such as RAID type or size—while their production volume and applications remain online.

Benefits of using LUN Migrator

The following benefits are attained when using CLARiiON LUN migration for moving data:

- Online and nondisruptive
- No impact on the surrounding infrastructure
- Transparent to VMware ESX/ESXi hosts, virtual machines, applications, and users
- Utilizes array resources
- Eliminates any reconfiguration pre- or post-migration

With traditional migration techniques, a number of manual tasks are required in addition to the copy operation. Even with VMware Storage vMotion, VMware administrators need to create and present the new storage device to the ESX/ESXi host, rescan for the new device, and create a VMFS datastore before they can migrate their data. CLARiiON LUN migration removes the requirement for any ESX/ESXi involvement or resources.

In contrast to Storage vMotion, however, there is no option for the granular movement of individual virtual disks on the datastore to different locations. All virtual disks (and anything else that may reside on that LUN) are migrated to the destination LUN.

When the migration operation completes, the original source LUN is destroyed and the new destination LUN assumes the **Nice Name**, **WWN**, and **LUN ID** of the source LUN. The ESX/ESXi Server always detects the same device because the migration is masked at the back end. Therefore, no other configuration changes are required anywhere in the environment.

Options available with LUN migration

LUN migration enables users to migrate LUN data from traditional RAID groups to storage pools. LUN migration also enables the migration of LUN data to:

- Different RAID types—for example, RAID 5 to RAID 1/0
- Different disk types—for example, FC to SATA to EFD
- Different LUN type—for example, thick to thin format
- More underlying spindles—that is, an eight-disk LUN versus a four-disk LUN
- A larger LUN—this requires the subsequent expansion of the VMFS datastore

Migration duration and performance

The duration of the migration depends on a number of factors. Some of these factors are user-configurable and some are not, while others depend solely on how busy the environment is at the time of migration.

For example, the size of the source LUN influences the duration of the migration but so too does the size of the destination LUN. The destination LUN can be larger in size so, in the case of a larger destination LUN, migration is a two-step process: copy and expansion. The expansion step requires additional time as the data is recalculated and relocated across the extra disk space. However, a regular migration to similar-sized LUNs is a simple, single-step copy operation.

Users can also configure the priority of the migration. There are four migration priorities: low, medium, high, and ASAP. These priorities set the rate of the migration and the subsequent utilization of SP resources. Users must decide between the impact the migration will have on production performance and how quickly they want the operation to complete.

Exceptions

The exceptions to using CLARiiON virtual LUN Migrator are:

- Private LUNs, such as reserved LUNs, clone private LUNs (CPLs), and write intent LUNs (WILs) in use by layered drivers such as EMC SnapView, SAN Copy, and MirrorView, cannot be migrated using the LUN migration tool.
- Individual component LUNs within metaLUNs are considered private and cannot be migrated independently. A metaLUN must be migrated as a whole object.
- LUNs supporting NFS-based storage cannot be migrated. Only block-level storage is supported.
- LUNs supporting NFS-based storage cannot be migrated (for example—when a CLARiiON storage system is providing the underlying the block storage for a Celerra's NFS file system). Only block-level storage is supported.
- Cross-array migrations are not supported. Operations are only supported within the same array.

CLARiiON LUN migration procedure

When the appropriate destination or target LUN is created on the CLARiiON, use the following steps in the Unisphere Management Console to start the migration process:

1. Select the source or original LUN to be migrated.
2. Right-click and select **Migrate**, as shown in Figure 47.




Figure 47: Selecting Migrate

3. From the dialog box displayed, select the target or destination LUN.
4. Select the **Migration Rate/Priority** (Low, Medium, or High).
5. Click **OK**.

The migration process begins automatically.

As the migration process is running, the system copies the contents and blocks of the original LUN to the destination LUN. The migration driver handles the synchronization of both LUNs, the switchover of host operations from the old LUN to the new LUN, and the subsequent removal and unbinding of the original LUN. The migration process occurs online during the host I/O with no interruption to production operations.

No other configuration changes are required, either before or after the migration completes, as the new LUN assumes all of the properties of the original LUN. From the host’s perspective, it is viewing the same LUN the entire time, without interruption.

The progress of the migration can be viewed on the LUN Migration Summary window, as shown in Figure 48.




Figure 48: Viewing the LUN Migration Summary window

All internal operations specific to the migration can be viewed within the CLARiiON Storage Processor Event Log, as shown in Figure 49.

102	2010-08-19	14:12:56	0x71570008	Migration Cancelled: 38. 00 00 04 00 02 00 58 00 d3 04 ... CKM00094100027	N/A	N/A	SPA
103	2010-08-19	14:12:56	0x71570003	Migration volume Destroyed: 38. 00 00 04 00 02 00 58 0... CKM00094100027	N/A	N/A	SPA
104	2010-08-19	14:12:56	0x7115000d	Internal Information only. PSM File "MigChkpt2" Deleted. ... CKM00094100027	N/A	N/A	SPA
105	2010-08-19	14:12:36	0x71570006	Migration Completed: 38. 00 00 04 00 02 00 58 00 d3 04... CKM00094100027	N/A	N/A	SPA
106	2010-08-19	14:12:07	0x71570006	Migration Completed: 36. 00 00 04 00 02 00 58 00 d3 04... CKM00094100027	N/A	N/A	SPA
107	2010-08-19	14:06:58	0x4600	'Start LUN Migration' called by 'Navi User Administrator' (1... CKM00094100027	N/A	N/A	SPA
108	2010-08-19	14:06:57	0x71570004	Migration Started: 38. 00 00 04 00 02 00 58 00 d3 04 00... CKM00094100027	N/A	N/A	SPA
109	2010-08-19	14:06:54	0x4600	'Start LUN Migration' called by 'Navi User Administrator' (1... CKM00094100027	N/A	N/A	SPA

Figure 49: Viewing the Event Log

Using LUN 38 as an example, it is clear when the migration process began and when it completed from the section of events shown in Figure 49. Figure 49 also shows that the system purged the original LUN upon completion of the migration.

When the migration process is complete, no further user intervention is required. The VMware administrator is not required to take any action at any stage as the migration operation is completely transparent.

**Scenario:
Migrating
applications
with CLARiiON
LUN Migrator**

To demonstrate how the CLARiiON LUN Migrator can be used to migrate applications, a mixed-application, 24x7 environment running virtualized Microsoft SQL 2005 and Exchange 2010 on a CLARiiON CX4-480 array was set up. The host environment consisted of Windows 2008 virtual machines running on VMware vSphere version 4.1 servers.

Under normal conditions (baseline), the SQL server, running the SQL DVD Store application, was running at approximately 22,000 commands per minute. The Exchange 2010 mail server was running 2,500 users.

The objectives of this migration operation were to demonstrate:

- The process of migration to new LUNs
- The duration of the migration operations
- The impact to application performance during the migrations

Table 8 provides a summary of the source and target storage.

Table 8: Source and target storage

	Source	Destination
SQL DVD store		
VMFS datastores	3 x 10 GB	3 x 10 GB
LUNs	3 x 11 GB	3 x 11 GB
RAID groups	2 x 4+4 RAID 1/0	2 x 4+4 RAID 1/0
Disks	16	16
Exchange 2010		
VMFS datastores	2 x 100 GB	2 x 100 GB
LUNs	2 x 100 GB	2 x 100 GB
RAID groups	2 x 4+1 RAID 5	2 x 4+1 RAID 5
Disks	10	10

The most important factors to observe are:

- The ease of use
- The associated impact of the migration was on the running applications
- Zero downtime or disruption

Environment topology

Figure 50 illustrates exactly where the migration process takes place. All the work is done on the back-end storage by the array.




Figure 50: Environment configuration

Array performance during the migration

In this scenario, all the resources required to complete the migration were managed by the CLARiiON array. By requiring the array to manage this operation, it was expected that an increased load would be visible on the array.

The performance chart in Figure 51 shows the duration of the Exchange 2010 data migration and the utilization percentage of the CLARiiON SPs during that time.




Figure 51: CLARiiON SP utilization during application migration

It can be seen that the application load prior to the migration had the array SPs running at approximately 6 percent utilization. When the migration was initiated, utilization increased to approximately 12–14 percent, before returning to 6 percent, once the migration process had completed.

The response times of the CLARiiON SPs also increased slightly during the migration, as shown in the performance chart in Figure 52.




Figure 52: CLARiiON SP response time

Application performance during the migration

To measure any potential impact on the application during the migration, Microsoft LoadGen was used to run a simulation of 2,500 Exchange 2010 users. The I/O profile of Exchange 2010 is such that the load on the back-end disks is quite light compared with previous versions of Exchange, so very little bottleneck was observed as the array and underlying disks were not being stressed.

Before any potential impact on the migration could be measured, a baseline run was measured, which consisted of a standard 2-hour test. A second test was then run that included the LUN migration. The LUN migration started 30 minutes into the 2-hour test run and completed in approximately 50 minutes.

The latencies gathered from Perfmon were identical for both test runs, as shown in Table 9.

Table 9: Exchange 2010 Perfmon latencies

Exchange 2010	Baseline	Including migration
Database seconds/Read	5 ms	5 ms
Database seconds/Write	1 ms	1 ms
Log seconds/Write	1 ms	1 ms

Similar behavior was observed during the SQL testing as shown in Table 10.

Table 10: SQL 2005 Perfmon latencies

SQL 2005	Baseline	Including migration
Operations per minute	21641	21322
Database seconds/Write	2 ms	2 ms
Log seconds/Write	1 ms	1 ms

In the case of the SQL test, a SQL DVD Store simulator was run for two hours, which under normal conditions operated at up to 21,500 operations per minute. When the migration was executed during the test run, the **Operations per minute** figure remained approximately the same, and the only impact on latencies was to the **Database seconds/Write** figure, which increased by one millisecond.

One of the main reasons that the applications experienced very little performance impact in this case was because the back-end resources had sufficient headroom to absorb the extra activity. For example, during the Exchange 2010 test run, the back-end disks were only 16 percent utilized under normal load conditions before the migration began. Therefore when the migration process required the disks to work harder (with extra reads and writes), the disks had plenty of cycles available. The same can be said of the CLARiiON SPs which, as shown previously in Figure 51, have plenty of room for increased activity.

If migrations were conducted with the array under a much heavier load, such as during peak activities in a day-to-day environment, and with system resources utilized more heavily, then it would be reasonable to expect that optimum application performance would be affected. For this reason, it is recommended that, as with any migration, such activity should be carefully planned and scheduled for non-peak intervals.

Conclusion

For users who want to migrate block-level data within a single storage array, CLARiiON virtual LUN Migrator is the easiest option for a number of reasons, including:

- Management is fully contained within Unisphere; the operation requires no reconfiguration to the overall SAN and requires the least amount of pre- or post-migration work.
 - CLARiiON virtual LUN Migrator functionality is native to CLARiiON FLARE and requires only that the storage administrator provision an appropriate target LUN before the migration can proceed.
 - No action is required by the VMware administrator. The migration process is completely transparent to ESX/ESXi servers and virtual machines. Applications and end users remain online throughout the process.
-

Data migration with VMware vCenter Site Recovery Manager

VMware vCenter SRM

This section discusses the use of VMware vCenter SRM as a migration tool to coordinate the migration of virtual machines and their data to another array or location.

VMware vCenter SRM is a disaster recovery management and automation solution for VMware Infrastructure. VMware vCenter SRM accelerates recovery by automating the recovery process and simplifying the management of disaster recovery plans.

VMware vCenter SRM and EMC RecoverPoint—using the RecoverPoint Storage Replication Adapter (SRA)—form a cooperative relationship to automate failover to a recovery site.

VMware vCenter SRM does not replicate any data; it coordinates the recovery process of the virtualized environment from one site to another. VMware vCenter SRM ensures that the server and virtual machine infrastructure is in place on the recovery site to restart services, once the replicated data is presented.

RecoverPoint does not provide for the recovery site infrastructure. It replicates the data from one site to another. RecoverPoint ensures the consistency of the data presented to the virtual infrastructure on the recovery site.

The RecoverPoint SRA for VMware vCenter SRM is a software package that enables VMware vCenter SRM to implement disaster recovery using RecoverPoint. RecoverPoint SRA supports VMware vCenter SRM functions, such as failover and failover testing, using RecoverPoint as the replication engine.

While many of the points apply equally, regardless of the migration technology used, this white paper deals with the use of Site Recovery Manager, leveraging RecoverPoint, for the underlying replication. The differences between this and other storage replication adapters are called out as appropriate. Storage replication adapters are available for IP Replicator, MirrorView, and SRDF.

Test environment

Figure 53 illustrates the architecture used for testing VMware vCenter SRM in this scenario.




Figure 53: Test environment for VMware vCenter Site Recovery Manager

Requirements for using VMware vCenter SRM as a migration tool

The architecture of VMware vCenter SRM requires a separate instance of VMware vCenter Server to be created for the remote site (and that at least one additional VMware vSphere host is present) so it can function correctly.

VMware vCenter SRM requires its own database and may be installed on the same server as VMware vCenter, although a separate server is recommended for performance. The RecoverPoint SRA must also be installed on whichever server is hosting VMware vCenter Site Recovery Manager. In this scenario, both VMware vCenter and Site Recovery Manager were hosted on a single virtual machine per site.

Any virtual machines migrated using VMware vCenter SRM, by definition, end up in a new vSphere environment.

The RecoverPoint SRA (as with other storage replication adapters) replicates at the block level, therefore a distinct set of target LUNs must be created on the remote site. Since everything on the source LUNs is replicated to a remote site, it is important to place all the elements of a virtual machine on these LUNs, including the **Guest OS Boot Volume**. Otherwise, the machine does not boot on the target site or the failover task may fail.

If memory reservations are being used on the source site, be aware that VMware vCenter SRM does not create memory reservations on the placeholder virtual machines. Therefore, additional space may be needed on the target datastore that hosts the VSWAP file for the failed-over virtual machine to boot. Alternatively, create a memory reservation on the placeholder virtual machine after the VMware vCenter SRM protection group has successfully created that remote placeholder virtual machine.

Using VMware vCenter SRM as a migration tool

The characteristics of using VMware vCenter SRM as a migration tool are:

- Migration of the virtual machines to the target vCenter instance is disruptive, in that the virtual machines must be powered down and then powered back up on the new site.
- Data can be synchronized in advance of the migration process so that the RTO is minimal.
- Migration is carried out at the LUN level, or a group of LUNs, depending on the configuration of the RecoverPoint consistency groups. Migration of virtual disks on an individual basis is not possible.
- All LUN operations, in terms of array failover operations, host rescans, addition of virtual machines to the inventory, and the powering on of those virtual machines is carried out automatically as part of the migration process, eliminating many of the pain points of manual migrations. The migration process can also be tested and validated prior to the event by means of the SRM Test Recovery Plan functionality.

Installing and configuring VMware vCenter SRM failovers

This paper does not provide detailed information about installing and configuring VMware vCenter SRM with RecoverPoint, as this is well documented in existing product guides.

The high-level stages are detailed in Table 11.

Table 11: High-level installation and configuration stages

Stage	Description
1	Install at least two vSphere hosts, one per site.
2	Install separate vCenter instances.
3	Configure local and remote storage for use as datastores or remote deployment managers (RDMS), as well as local and remote RecoverPoint journals, as appropriate.
4	Configure and synchronize RecoverPoint consistency groups for all relevant LUNs.
5	Install VMware vCenter SRM.
6	Install the relevant storage replication adapter.

When using RecoverPoint as the replication technology, it is important to ensure that the RecoverPoint consistency groups are configured to be managed by VMware vCenter Site Recovery Manager. This is done in RecoverPoint at the **Consistency Groups** level, as shown in Figure 54.




Figure 54: Configuring consistency groups

Once the VMware vCenter SRM instance is up and running, the process of migrating virtual machines to another set of storage and vSphere hosts follows the standard procedure for creating protection groups and recovery plans.

From the VMware vCenter SRM plug-in in the vCenter client, click **Create Protection Group** and name the group, as shown in Figure 55.




Figure 55: Configuring the Protection Group Name

The **Datastore Group** list detected by the SRA is presented. Select one or more groups to manage at the same time. A list of the affected virtual machines is displayed, as shown in Figure 56.




Figure 56: Managing a Datastore Group

The migration process is as simple as executing a standard VMware vCenter SRM failover.

Once the migration is complete, the original site still retains a powered-off virtual machine, and the source storage array or LUN still contains a copy of the data. Appropriate housekeeping needs to be carried out on the source.

Impact on production performance

Apart from the obvious impact of downtime during the migration itself, it is also important to consider the effect on the production systems while continuous data replication is being carried out by VMware vCenter Site Recovery Manager, and the underlying RecoverPoint SRA and its associated technology.

The impact on production depends on several factors, including:

- The distance and latency (if any) between the two storage arrays, if a cross array migration is being performed
- The configuration of the local and remote LUNs, and their ability to sustain the additional read I/O involved in the initial synchronization
- The method of synchronization used (synchronous or asynchronous)

In this test, two separate arrays were used, at zero distance. The impact of the initial data synchronization was recorded by measuring the response time of the Exchange database and log devices across a 2-hour LoadGen test. The initial data replication began approximately 45-50 minutes into the test, with the following results as shown in Figure 57 and Figure 58.




Figure 57: Average response time of the Exchange database

In this case, there is no mistaking the period during which the initial synchronization took place. A very distinctive increase in response time was observed in the disk response times, as recorded by the Exchange host.




Figure 58: Average response time of the log devices

The response times were still well within normal guidelines, so users may not notice any discernable difference but, as with any migration, this should be factored into the scheduling of such tasks in a real environment.

Using RecoverPoint SRA for local array or LUN migration

RecoverPoint provides the option to replicate to a separate storage array or back up to the original source array. The RecoverPoint appliances and cluster simply replicate data at a block level between source and target LUNs, regardless of where they reside on the SAN.

This functionality is unique with respect to the EMC storage replication adapters, since all the others are intended for use with remote replication technologies that require the target LUN to be on a separate array. However, it is only viable in scenarios using a host-based or SAN-based RecoverPoint splitter and does not work with the CLARiiON CX splitter.

A RecoverPoint splitter can be host-based, SAN-based, or array-based (CLARiiON only). It is responsible for splitting or duplicating all write I/O destined for a RecoverPoint-protected LUN, and sending copies to both the original target and to a RecoverPoint appliance, which then sends the duplicate I/O to a journal LUN and one or more target replica volumes.

When using VMware vCenter SRM as a migration tool, rather than its traditional use as a disaster recovery tool, the term *site* can therefore be interpreted as either a geographically separate location (traditional VMware vCenter SRM usage), or as a separate vCenter instance in the same location.

Because of this, it is technically possible to use VMware vCenter SRM and RecoverPoint as a method of migrating virtual machines from one array to another on a local site, as part of an array migration or as an upgrade. This has its advantages over SAN Copy in that VMware vCenter SRM coordinates all the steps that would ordinarily be required after a SAN Copy migration. These tasks include:

- Rescanning the target vSphere hosts
- Renaming and discovering datastores, if required
- Adding virtual machines to the inventory
- Completing any required virtual machine customization through Inventory Mapping and individual customizations
- Powering off all failed-over virtual machines in user-specified categories of importance

Conclusion

VMware vCenter SRM, in combination with RecoverPoint, provides a simple and effective method of carrying out inter-site virtual machine migrations. It enables replication of the datastores and RDMS that make up those virtual machines, in advance of the need to failover. It also coordinates all the tasks required to make these virtual machines operational on the target site.

In addition, it can be used as a technique to migrate virtual machines to a new array or to different LUNs on the same site, when using the RecoverPoint replication option. This may provide value in situations where techniques such as LUN migration or Storage vMotion are not appropriate or available.

Summary of data migration techniques

Comparing the data migration techniques

Table 12 provides a comparison of the data migration techniques described in this white paper.

Table 12: Data migration techniques

	VMware Converter	VMware Storage vMotion	CLARiiON LUN migration	CLARiiON SAN Copy	VMware vCenter SRM	EMC VPLEX Metro
Migration granularity level	Virtual machine	Virtual disk	Array LUN	Array LUN	Consistency Groups defined by RecoverPoint Storage Replication Adapter	Array LUN
Storage type	Heterogeneous	Heterogeneous	Homogenous	Heterogeneous (minimum of one CX array)	Heterogeneous	Heterogeneous
Disruption level	Disruptive	Nondisruptive	Nondisruptive	Disruptive	Disruptive	Nondisruptive
Storage migration type	In- or cross-array migration	In- or cross-array migration	In-array migration	In- or cross-array migration	Cross-array migration*	In- or cross-array migration
Manual steps post-migration	Required	Not required	Not required	Required	Not required	Not required
Migration between thick and thin devices	At vSphere level	At vSphere level	At array level	At array level	At array level	At array level
Physical or virtual environments	Both (can convert physical to virtual)	Virtual only	Both	Both	Virtual only	Both

* If used with RecoverPoint and either the SAN or host-based splitter, then in-array migration is also possible.

Scenario 1: Migration of individual virtual disks

In general, the task of migrating individual virtual disks is undertaken to redistribute data to datastores with different performance characteristics, or to datastores that are stored on LUNs of a different RAID type.

VMware Storage vMotion provides the only effective solution for this move. It provides the major benefit of enabling individual virtual disks to be moved to any datastore that can be accessed by the VMware vSphere server hosting the virtual machine.

As long as that storage can be presented to that host, then Storage vMotion can move it online, and without disruption to production. There may be an impact on the performance of the virtual machine, depending on the configuration of the source and target LUNs, as the Storage vMotion operation places additional I/O overhead on both the target and source LUNs.

The effect of this additional overhead may potentially be reduced if you are utilizing vSphere 4.1 and have an array that can take advantage of the VMware vSphere VAAI Full Copy feature, which can offload a significant amount of work to the storage array.

**Scenario 2:
Migration of
datastore
contents to
another LUN on
same array**

For the migration of datastore contents to another LUN on same array, a number of techniques can be deployed, depending on the purpose of this migration. For example, VMware Storage vMotion can be used to migrate the virtual disks, one virtual machine at a time, to the new datastore, assuming that a new target datastore had been created and assigned to the VMware vSphere cluster.

If the intention is to move all the content from the source datastore to a previously existing datastore, then Storage vMotion is again the only viable option.

If, however, this is not the case and the intention is to move the datastore to a new LUN with different characteristics (for example, RAID type, performance, thin versus thick, and so on), then a number of other options present themselves.

SAN Copy is capable of performing these operations but requires the virtual machines to be quiesced prior to transferring operations to the new LUN.

Both CLARiiON LUN Migrator and VPLEX Metro are also capable of transparently migrating the data, using back-end storage processes, to the VMware vSphere host, and without interruption of service to the production virtual machines. This is assuming that the source LUNs are already present on a CLARiiON or VPLEX system respectively, before the migration.

**Scenario 3:
Migration of
entire contents
of a datastore
to a LUN on
another array**

There are a number of reasons for migrating a datastore to a new array: better performance or functionality, balancing storage usage, or simply because an older array is being decommissioned.

VMware Storage vMotion is a potential candidate if the new source array can be presented to the original VMware vSphere host(s). This is ideal when a nondisruptive migration is required. However, it requires each virtual machine to be migrated individually, which may be time consuming.

SAN Copy can be used to migrate data between arrays, as long as either the source or target array supports SAN Copy. Advantages include the ability to move the contents of a single datastore (and not necessarily all of the virtual machines) and the ability to maintain the virtual machine on the same VMware vCenter instance. SAN Copy can also be used to migrate to a new VMware vCenter instance, using new storage, by presenting the SAN Copy target(s) to the new instance, and adding the virtual machines to the inventory.

A disadvantage is that a number of manual steps are needed to remove the original LUNs and add the new LUNs to the VMware vSphere cluster, and the associated downtime of the virtual machines involved in carrying out these tasks.

VMware vCenter SRM can be used in this scenario, leveraging a replication technology such as RecoverPoint (making it array-agnostic) or another storage replication adapter, specific to the array type being used.

The advantage of this approach is that SRM automates all the manual tasks normally required to complete a SAN Copy migration.

The disadvantages are the need for a separate VMware vCenter instance, the vSphere host could be moved from original site, and the cost of VMware vCenter SRM. It also requires a period of downtime for the virtual machine during the transfer to the new array or vSphere host and the entire virtual machine must be moved to the new array.

VPLEX Metro can be used in this case, assuming that the source LUN is already encapsulated by a VPLEX appliance. In this scenario the data can be migrated transparently to the hosts or virtual machines.

**Scenario 4:
Migration of the
entire virtual
infrastructure
to another
geographical
location**

Technically, several options can be used for the migration of an entire virtual infrastructure to another geographical location.

VMware vMotion over synchronous distances is now a valid option with VPLEX, assuming the source LUN is already encapsulated by VPLEX. This provides the only nondisruptive solution in this scenario

VMware vCenter SRM is a natural fit for this scenario. It removes all the manual steps involved in both VMware Storage vMotion and SAN Copy, and coordinates the entire failover and migration to the remote site.

If it is possible to present storage from the remote site to the local VMware vSphere host, Storage vMotion can be used. However, this requires the following conditions:

- Removal of the virtual machines from the inventory on the source site
- Removal of the datastores from the local VMware vSphere host
- Addition of datastores to the remote VMware vSphere host
- Addition of virtual machines to the inventory on the remote VMware vSphere host
- Ensuring that the VMX configuration of the virtual machines is correct
- Powering on the virtual machines

SAN Copy can provide similar functionality to the Storage vMotion option with the following exceptions:

- It does not require each virtual machine to be migrated individually.
- It does not require the remote storage to be presented to the local VMware vSphere cluster.
- It does not require the creation of an additional datastore, as the datastore is copied at a storage level to the remote site.

**Scenario 5:
Migration from
physical to
virtual or inter-
hypervisor**

The only solution in a migration from a physical to a virtual or inter-hypervisor scenario is VMware vCenter Converter, which is available in two different versions: VMware vCenter Converter Standalone and VMware vCenter Converter, which is integrated with the VMware vCenter Server.

Both versions enable:

- Online hot-cloning of a physical or virtual machine to a VMware virtual machine.
 - Incremental cloning of a hot-cloned machine before transferring operations to the newly commissioned virtual machine.
 - Cold cloning of a physical or virtual machine, where that machine is booted from a VMware vCenter Converter Boot CD, but where the actual OS or data to be migrated is not actively running production.
-

Conclusion

Summary

Data center migration is a complex and often risk-heavy undertaking. Organizations must carefully plan and execute each step of the process. Migration windows are short and schedule delays are costly. Operations cannot stop running while the data center is being moved and the data is continuously changing during the planning and migration process. The migration process must therefore be flexible enough to accommodate changes as they occur.

Findings

This comparative study highlights several approaches to data migration for VMware vSphere and provides valuable insight for customers planning a data center migration of their virtual information infrastructure environments. There is no single solution to all data migrations. Each type of data migration must be considered in its own right and due consideration given to the conditions and components involved. Some scenarios will lend themselves to certain methods of data migration whereby customers are able to fully realize data migration flexibility at the server and storage layers of the environment.

Next steps

To learn more about this and other solutions contact an EMC representative or visit: www.emc.com.

References

White papers

For additional information, see the white papers listed below.

- *Using VMware Virtualization Platforms with EMC VPLEX—Best Practices Planning*
 - *EMC SAN Copy—A Detailed Review*
 - *EMC CLARiiON Integration with VMware ESX—Applied Technology*
 - *EMC CLARiiON Best Practices for Performance and Availability: Release 30.0 Firmware Update*
-

Other documentation

For additional information, see the documents listed below.

- *VMware ESX Configuration Guide*
 - *VMware vCenter Site Recovery Manager Administration Guide*
 - *VMware Converter Standalone 4.3 User's Guide*
 - *vSphere Datacenter Administration Guide*
-