

EMC Symmetrix VMAX and VMware Virtual Infrastructure

Applied Technology

Abstract

This white paper describes EMC[®] Symmetrix[®] VMAX[™] system features and functionality relevant to VMware Virtual Infrastructure environments. New Enginuity[™] features, such as Auto-provisioning Groups and Enhanced Virtual LUN Technology, and their use in VMware Virtual Infrastructure environments, are discussed. The paper also presents enhancements in EMC Solutions Enabler and new tools in support of VMware data center products.

May 2010

Copyright © 2009, 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 h6209.1

Table of Contents

Table of Figures	4
Executive summary	5
Introduction	5
Audience	5
Overview	6
Symmetrix VMAX	6
Connectivity	8
Multipathing and static load balancing	10
EMC Storage Viewer for Virtual Infrastructure Client	12
Engenuity release 5874 and maximum device size	14
Auto-provisioning Groups	19
Provisioning storage to a VMware Virtual Infrastructure	20
Modifying existing views	23
Use Case: Expanding VMware cluster groups by sharing datastores	24
Use Case: Providing additional capacity to an existing VMware cluster group	26
Scalable and flexible design of Auto-provisioning Groups	26
Hybrid compute environments with boot from SAN for blade servers	27
Isolation of workloads to a subset of the compute environment	31
Solutions Enabler in VMware Virtual Infrastructure environments	34
Virtual LUN technology	36
Rules regarding Virtual LUNs	37
Migrating VMFS datastores	37
Use Case: Migrating a VMFS volume from RAID 5 to RAID 1	37
Virtual Provisioning enhancements	41
SRDF enhancements	42
SRDF/Extended Distance Protection (SRDF/EDP)	42
RDF Engenuity Consistency Assist (RDF-ECA)	42
SRDF/A adding or removing devices	42
Remote TimeFinder/Clone to SRDF R1 restore	43
SRDF/Star with the R22 device	43
TimeFinder enhancements	43
Cascaded TimeFinder/Clone	43
TimeFinder/Snap Recreate	43
Open Replicator enhancements	43
Cold push from TimeFinder/Snap snapshot (VDEV)	43
Conclusion	44
References	44

Table of Figures

Figure 1. Symmetrix VMAX system features.....	6
Figure 2. VMAX Engine	7
Figure 3. Front-end I/O placement	7
Figure 4. Connecting ESX Servers to a single engine Symmetrix VMAX.....	9
Figure 5. Connecting ESX Servers to a Symmetrix VMAX with more than one engine	9
Figure 6. Datastore Paths using two HBAs and two directors	10
Figure 7. Setting a preferred path using VMware Virtual Infrastructure Client	11
Figure 8. Configuring datastores on an ESX Server with alternating paths.....	12
Figure 9. Datastore subview of the ESX Server context view of EMC Storage Viewer	13
Figure 10. LUN subview of the ESX Server context view of EMC Storage Viewer	13
Figure 11. Virtual Machine context view of EMC Storage Viewer	14
Figure 12. Symmetrix disk groups seen through Symmetrix Management Console.....	15
Figure 13. Creating a 240 GB Symmetrix device with EMC Solutions Enabler	16
Figure 14. RAID 5 disk consumption for a 240 GB volume in Symmetrix Management Console	17
Figure 15. View of a disk group containing 1 TB 7200 rpm SATA drives using SMC	18
Figure 16. Creating a 240 GB RAID 1 volume with EMC Solutions Enabler.....	18
Figure 17. Disk consumption of a 240 GB RAID 1 volume over two 1 TB SATA drives.....	19
Figure 18. Adding devices to the storage group	21
Figure 19. Adding ports to the port group	21
Figure 20. Adding initiators to initiator groups	22
Figure 21. Creating a view	22
Figure 22. Performing a rescan of storage adapters on a VMware ESX Server	23
Figure 23. Newly discovered Symmetrix devices shown using EMC Storage Viewer for Virtual Infrastructure Client	23
Figure 24. Adding initiators to a view	24
Figure 25. Detailed contents of a view	25
Figure 26. Expanding the VMware cluster by sharing a VMFS datastore.....	26
Figure 27. Boot from SAN information using EMC Storage Viewer for Virtual Infrastructure Client.....	27
Figure 28. Properties of a view through Symmetrix Management Console.....	28
Figure 29. Cascading initiator groups in Solutions Enabler.....	29
Figure 30. Creating and displaying a view with cascaded initiator groups.....	30
Figure 31. Displaying the details of a view for a cascaded initiator group	30
Figure 32. Shared Symmetrix VMAX devices through Auto-provisioning Groups.....	31
Figure 33. Display of VMware Virtual Infrastructure supporting a varied workload	32
Figure 34. Display showing the hosts connected to production and QA datastores	32
Figure 35. The view for presenting the QA_Datastore devices to the QA cluster.....	33
Figure 36. The view for presenting the production datastore to the VI environment.....	34
Figure 37. Adding credentials to the SYMAPI database for ESX Server or vCenter Server.....	35
Figure 38. Using the sympd command to display virtual disk mapping information	36
Figure 39. Symmetrix device before migration using EMC Storage Viewer for Virtual Infrastructure Client.....	37
Figure 40. Detailed Symmetrix device information before migration	38
Figure 41. Validating a VLUN migration	38
Figure 42. Establishing a VLUN migration session	39
Figure 43. Querying a VLUN migration session during and after migration	40
Figure 44. Verifying a VLUN migration session during and after migration	40
Figure 45. Detailed Symmetrix device information after migration.....	41
Figure 46. Symmetrix device information after migration using EMC Storage Viewer for VMware Infrastructure Client	41

Executive summary

The EMC® Symmetrix® VMAX™ Series with Enginuity™ provides an extensive offering of new features and functionality for the next era of high-availability virtual data centers. With advanced levels of data protection and replication, the Symmetrix VMAX is at the forefront of enterprise storage area network (SAN) technology. Additionally, the Symmetrix VMAX array is faster, larger, and more efficient, which allows it to transparently optimize service levels without compromising the capability to deliver performance on demand. These capabilities are of great value for large virtualized server farms such as VMware Virtual Infrastructure.

The expansion of VMware Virtual Infrastructure implementations in current data centers drives the need for larger consolidated storage arrays. As demand increases, storage administrators will need the usual tasks involved with managing storage to be consolidated, simplified and consume less time to complete. One task in the process of managing the storage environment is the provisioning of storage. The redundant and common steps associated with this process are simplified in Symmetrix VMAX with the introduction of a new method for device masking and mapping called Auto-provisioning Groups. This feature greatly reduces the time it takes to map and mask Symmetrix devices and present them to a VMware Virtual Infrastructure. The time saved with Auto-provisioning Groups is magnified when allocating large amounts of shared storage to large virtual server farms, thus allowing more time for storage administrators to focus on more important, critical tasks. In addition to Auto-provisioning Groups, Symmetrix VMAX provides support for larger Symmetrix devices and automatic creation of metavolumes. These enhancements reduce the instances where the storage administrator is forced to perform additional activities when provisioning storage to a VMware Virtual Infrastructure environment.

The value of business data normally changes with time. Thus the importance of the virtual machines utilizing the data also changes with time. The VMware Virtual Infrastructure provides users with mature tools to manage the lifecycle of VMware virtual machines. Symmetrix VMAX introduces the ability to nondisruptively and transparently move data volumes within the storage array, empowering the storage administrator to provide lifecycle management for business data. The powerful new feature called Virtual LUN migration allows Symmetrix devices to be moved between RAID levels or disk types with no disruption to the VMware Virtual Infrastructure accessing the device. This feature allows the Symmetrix VMAX system to support a wide range of service levels within a single storage array and provides a key building block for implementing Information Lifecycle Management.

Introduction

The Symmetrix VMAX and VMAX SE systems running the Enginuity operating environment are a new entry in the enterprise-class Symmetrix storage array offerings from EMC. Symmetrix VMAX is a high-end, scalable storage array that provides an extensive offering of new features and functionality for the next era of high-availability virtual data centers. The array also offers advanced levels of data protection and replication, placing it at the forefront of enterprise SAN technology. Additionally, the Symmetrix VMAX is more energy efficient, faster, and larger and can transparently optimize service levels without compromising the capability to deliver performance on demand that is required for large virtualized server farms such as VMware Virtual Infrastructure.

This white paper discusses new features introduced with the release of Symmetrix VMAX that are the most relevant in a VMware Virtual Infrastructure deployment, including Auto-provisioning Groups and Virtual LUN migration technology. The white paper also discusses best practices when connecting a VMware Virtual Infrastructure environment to Symmetrix VMAX.

Audience

The audience for this white paper is VMware administrators and storage administrators responsible for managing and supporting VMware Virtual Infrastructures utilizing the Symmetrix VMAX storage array.

Overview

Symmetrix VMAX

The Symmetrix VMAX system with the Engenuity operating environment is a new enterprise-class storage array that is built on the strategy of simple, intelligent, modular storage. The array incorporates a new scalable fabric interconnect design that allows the storage array to provide improved performance and scalability for demanding enterprise storage environments while maintaining support for EMC's broad portfolio of platform software offerings. The storage array seamlessly grows from an entry-level configuration with a single, highly available VMAX Engine and one storage bay into the world's largest storage system with eight engines and 10 storage bays. The largest supported configuration is shown in Figure 1. The figure also shows the range of configurations that is supported by the Symmetrix VMAX storage array.

- 2-16 director boards
- 40-2,400 disk drives
- Up to 2 PB usable capacity
- Up to 128 Fibre Channel ports
- Up to 64 FICON ports
- Up to 64 Gig-E/iSCSI ports
- Up to 1 TB global memory




Figure 1. Symmetrix VMAX system features

The enterprise-class applications in a modern data center are expected to be always available. The design of the Symmetrix VMAX storage array enables it to meet this stringent requirement. The hardware and software architecture of the Symmetrix VMAX storage array allows capacity and performance upgrades to be performed online with no impact to production applications. In fact, all configuration changes, hardware and software updates, and service procedures are designed to be performed online and nondisruptively. This ensures that customers can consolidate without compromising availability, performance, and functionality, while leveraging true pay-as-you-grow economics for high-growth storage environments.

Symmetrix VMAX can include two to 16 directors inside one to eight VMAX Engines. Each VMAX Engine has its own redundant power supplies, cooling fans, SPS Modules, and Environmental Modules. Furthermore, the connectivity between the Symmetrix VMAX array engines provides a direct connection from each director to every other director, creating a Virtual Matrix. Each VMAX Engine has two directors

that can offer up to eight ports each, therefore allowing up to 16 ports per VMAX Engine. Figure 2 shows the typical setup of a Fibre Channel-only VMAX Engine.




Figure 2. VMAX Engine

In Symmetrix storage arrays, directors were set up differently; front-end directors, back-end directors, and memory boards were separate entities. In Symmetrix VMAX these have been combined into integrated directors to increase the number of possible directors in the Symmetrix. Each integrated director has three parts: the back-end director, the front-end director, and the cache memory module. The back-end director consists of two back-end I/O modules with four logical directors A, B, C, and D (not labeled in the figure) that connect directly into the integrated director. The front-end director consists of two front-end I/O modules with four logical directors (labeled E, F, G, and H) that are located in the corresponding I/O annex slots. The front-end I/O modules are connected to the director via the midplane (Figure 3).




Figure 3. Front-end I/O placement

The cache memory modules are located within each integrated director, each with eight available memory slots. Memory cards range from 2 to 8 GB, consequently allowing anywhere between 16 to 64 GB per integrated director. It is important to note that mixing memory card sizes within a director is not allowed; they must all be 2 GB, 4 GB, or 8 GB in size. For added redundancy, Symmetrix VMAX employs the use of mirrored cache, so memory is mirrored across engines in a multi-engine setup. In case of a single-engine configuration, the memory is mirrored inside the engine over the two integrated directors.

Symmetrix VMAX supports the following drive types:

Table 1. Disk drive support

Drive type	Rotational speed	Capacity
4 Gb/s FC	15k	146 GB, 300 GB, 450 GB
4 Gb/s FC	10k	400 GB
SATA	7.2k	1 TB
4 Gb/s Flash (SSD)	n/a	200 GB, 400 GB

A more detailed discussion of the Symmetrix VMAX storage array can be found in the product guide available on Powerlink[®], EMC's customer- and partner-only extranet.

Connectivity

The changes in the Symmetrix VMAX architecture and the Enginuity operating environment have greatly increased its performance, reliability, and availability. However, as a result of architectural changes there are new recommendations for connecting VMware ESX Servers to the array. The new recommendations ensure the highest level of connectivity and availability to the VMware Virtual Infrastructure even during abnormal operations.

Each VMware ESX Server in the VMware Virtual Infrastructure environment should have at least two physical HBAs, and each HBA should be connected to at least two different front-end ports on different directors. This configuration ensures continued use of all HBAs on the VMware ESX Server even if a front-end port goes down for maintenance activities on the director. Furthermore, the additional connectivity avoids unnecessarily idle equipment and reduces the performance impact caused by an offline director.

If the Symmetrix VMAX in use has only one VMAX Engine then each HBA should be connected to the odd and even directors within it. Connectivity to the Symmetrix front-end ports should consist of first connecting unique hosts to port 0 of the front-end directors before connecting additional hosts to port 1 of the same director and processor (Figure 4).




Figure 4. Connecting ESX Servers to a single engine Symmetrix VMAX

If multiple VMAX Engines are available, the HBAs from the VMware ESX Servers in the VMware Virtual Infrastructure environment should be connected to different directors on different VMAX Engines. In this situation, the first VMware ESX Server would be connected to four different processors on four different directors over two VMAX Engines instead of four different processors on two directors in one VMAX Engine (Figure 5).




Figure 5. Connecting ESX Servers to a Symmetrix VMAX with more than one engine

When the VMware ESX Server is connected to the Symmetrix VMAX array using the best practice recommendation discussed here, the VMware kernel will associate four paths to each device presented from the array. An example of this is shown in Figure 6 in which a datastore called "Test_Datastore" has been created on a device that the VMware kernel can access using one of the four possible paths. The example depicts a VMware ESX Server connected to a Symmetrix array that only has one VMAX Engine and therefore its connectivity is limited to two directors.




Figure 6. Datastore Paths using two HBAs and two directors

As more VMAX Engines become available, the connectivity can be scaled as needed. These methodologies for connectivity ensures all front-end directors and processors are utilized, providing maximum potential performance and load balancing for VMware Virtual Infrastructure environments connected to the Symmetrix VMAX storage arrays.

Multipathing and static load balancing

VMware ESX Server provides native channel failover capabilities. ESX Server, by default, assigns the first discovered path to an EMC Symmetrix disk as the preferred path with a “fixed” failover policy. Furthermore, this path is also configured as the active path and used for sending I/O to that disk. All other paths are used as a passive failover path and utilized only if the active path fails. In such configurations, all I/Os from the ESX Server queues on the first available HBA in the system, while the other HBAs are not actively used until a failure on the primary HBA is detected. This behavior leads to an unbalanced configuration on the ESX Server and on the Symmetrix VMAX storage system.

At this time, the best means of load balancing and multipathing on an ESX Server is through the static assignment of alternating preferred paths to Symmetrix VMAX devices. Given the recommendations in the “Connectivity” section on page 8, every ESX Server should have, at the very least, four separate paths to the array. With this approach, VMFS datastores residing on Symmetrix VMAX should have their preferred paths balanced evenly over the available paths.

As shown in Figure 7, setting or changing the preferred path can be done using the VMware Virtual Infrastructure Client. In Figure 8 there are four datastores, each linked to front-end ports 7E:0, 7G:0, 8E:0, and 8G:0 on the Symmetrix VMAX array. To maximize throughput and avoid wastefully idle equipment each datastore has its preferred path set to a different physical path.




Figure 7. Setting a preferred path using VMware Virtual Infrastructure Client




Figure 8. Configuring datastores on an ESX Server with alternating paths

These recommendations will provide greater resiliency in a VMware ESX Server environment by reducing the performance impact from physical hardware failures or from intentionally bringing equipment offline. Furthermore, presenting a device on multiple paths allows for better load balancing and reduced sensitivity to storage port queuing. This is extremely important in environments that share Symmetrix VMAX storage array ports between ESX Servers and other operating systems.

EMC Storage Viewer for Virtual Infrastructure Client

VMware Infrastructure is the most widely deployed software suite for optimizing and managing IT environments through virtualization. With the maturing of this technology as a platform of choice for deploying mission-critical x86 applications, business activities are frequently completed from within the VMware Virtual Infrastructure Client. To aid in this goal, EMC now provides EMC Storage Viewer for Virtual Infrastructure Client. This tool adds additional capabilities to the Virtual Infrastructure Client, so that users may now view detailed storage specific information, previously impossible to obtain without using one or two separate applications outside of Virtual Infrastructure Client. EMC Storage Viewer provides simple, read-only storage mapping functionality for various EMC storage-related entities that exist within Virtual Infrastructure Client, including datastores, LUNs, and SCSI targets. The storage information displayed through EMC Storage Viewer allows distinction among the types of storage used, the specific arrays and devices presented, the paths that are used for the storage, and individual characteristics of the existing storage.

EMC Storage Viewer provides two main views—the ESX Server context view and the Virtual Machine context view. The datastore subview of the ESX Server context view (Figure 9) provides detailed information about EMC storage associated with VMFS datastores visible on the selected ESX host (shown in Figure 9).




Figure 9. Datastore subview of the ESX Server context view of EMC Storage Viewer

The EMC Storage Viewer can also display detailed information about EMC storage visible to the ESX Server in the LUN subview. An example of this is shown in Figure 10. The information displayed includes protection type, serial number of the array, Engenuity level, device type, metavolume information, whether it's a virtually provisioned volume, and more.




Figure 10. LUN subview of the ESX Server context view of EMC Storage Viewer

The Targets subview drills down to all of the SCSI targets and displays detailed information on each, allowing for the most efficient way ever to map specific Symmetrix arrays to their respective SCSI target numbers on the ESX Server.

The second view is the Virtual Machine context view (Figure 11). This view displays detailed information about EMC storage devices associated with the selected virtual machine. The Symmetrix device information is provided for all of the datastores and/or raw device mappings in use by the virtual machine.




Figure 11. Virtual Machine context view of EMC Storage Viewer

The Virtual Machine context view provides similar information to the ESX Server context view, but instead limits the information to only those EMC storage devices associated with that virtual machine.

Virtual Infrastructure Client provides an inadequate amount of detail for mapping data sources presented to the ESX Server or virtual machines to their underlying storage devices. EMC Storage Viewer for Virtual Infrastructure Client fills this gap. EMC Storage Viewer allows users to view and decipher all storage entities in ways that the Virtual Infrastructure Client has never allowed before. Due to the combination of the EMC Storage Viewer and VMware's Virtual Infrastructure Client, managing a virtual environment attached to EMC storage is easier and far more efficient than ever before. For these reasons, it is highly recommended that EMC Storage Viewer be installed to manage any VMware Virtual Infrastructure environment making use of the Symmetrix VMAX array. An in-depth description of the functionality provided by EMC Storage Viewer for VMware Virtual Infrastructure is available in the white paper *Using EMC Storage Viewer for Virtual Infrastructure Client* on Powerlink.

Engenuity release 5874 and maximum device size

VMware Virtual Infrastructure abstracts a virtual machine to a collection of files that are stored in a datastore. A datastore is either a VMware file system created on a Fibre Channel or iSCSI attached device, or a NFS file system. The VMware datastore is treated by VMware administrators as a pool of storage and, therefore, can range from several hundred gigabytes to terabytes in size.

The Symmetrix DMX™ storage arrays running Engenuity release 5773 supported creation of a single Symmetrix device no larger than 60 GB. The need for larger LUN sizes for VMware Virtual Infrastructure was accommodated by creating an aggregation of individual Symmetrix devices called metavolumes. In addition to the support for metavolumes, the Symmetrix VMAX storage arrays and Engenuity release 5874 support creation of a 240 GB Symmetrix device. It is tempting to use the larger devices instead of metavolumes for VMware Virtual Infrastructure. However, although this may be acceptable for most of the VMware Virtual Infrastructure environments, as discussed in the following sections it may not always be the most appropriate choice.

Figure 12 shows the different disk groups available on a Symmetrix VMAX storage array. Also shown in this figure are the details for disk group 2, which is a collection of 6,146 GB, 15,000 rpm drives.

The screenshot displays the EMC Symmetrix Management Console interface. On the left, a tree view shows the hierarchy of storage components, with 'Disk Groups (4)' selected and '2' highlighted. The main pane shows the 'Properties - 000194900080 Local/2' window. The 'Data Disks (6)' tab is active, displaying a table of disk details. A red box highlights the first six rows of this table. Below, the 'General' tab shows various disk attributes, with 'Disk RPM' and 'Disk Capacity (GB/Cyl)' circled in red.

Disk	Int	TID	Hypers	Group Num	Cap (GB/Cyl)	Free C
DF-7C:D 7	D	7	4	2	136 / 146014	136 / 145862
DF-7D:C 7	C	7	4	2	136 / 146014	132 / 142230
DF-7D:D 8	D	8	4	2	136 / 146014	136 / 145862
DF-8C:C 7	C	7	3	2	136 / 146014	132 / 142239
DF-8C:D 8	D	8	4	2	136 / 146014	132 / 142230
DF-8D:D 7	D	7	3	2	136 / 146014	132 / 142239

Property	Value
Director	DF-7C
Interface	D
Target ID	7
Disk Group Number	2
Vendor ID	SEAGATE
Product ID	SXT14685FC
Product Revision	T146155
Serial ID	3LN667C2
Disk RPM	15000
Hypers	4
Block Size	520
Disk Capacity (GB/Cyl)	136 / 146014
Free Capacity (GB/Cyl)	136 / 145862
Actual Capacity (GB/Cyl)	136 / 146014

Figure 12. Symmetrix disk groups seen through Symmetrix Management Console

As shown in Figure 13, the attempt to create a RAID 1 protected, 240 GB Symmetrix device on this disk group fails. The Enginuity operating environment attempts to find a contiguous set of blocks on a single physical drive in the disk group to accommodate the request. The code does not attempt to automatically spread the request across all available physical disks in the disk group. Since it is impossible to create a 240 GB device on a 146 GB hard disk drive, the request fails.

```
root@licoe069:~
[root@licoe069 root]# cat /tmp/create_dev_mir
create dev count=1, size=262668 cyl, emulation=fba, config=2-Way-Mir,
disk_group=2, mvs_ssid=A;
[root@licoe069 root]#
[root@licoe069 root]# symconfigure -sid 194900080 -file /tmp/create_dev_mir -nop commit

A Configuration Change operation is in progress. Please wait...

Establishing a configuration change session.....Established.
Processing symmetrix 000194900080
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Initiating COMMIT of configuration changes.....Queued.
COMMIT requesting required resources.....Failed.
Definition 1 is in error:
Could not find disks that satisfy our mirror/RAID policy

Closing configuration change request.....Closed.
Terminating the configuration change session.....Done.

The configuration change session has failed.

[root@licoe069 root]# symconfigure -sid 194900080 -file /tmp/create_dev_raid5 -nop commit

A Configuration Change operation is in progress. Please wait...

Establishing a configuration change session.....Established.
Processing symmetrix 000194900080
Performing Access checks.....Allowed.
Checking Device Reservations.....Allowed.
Initiating COMMIT of configuration changes.....Queued.
COMMIT requesting required resources.....Obtained.
Step 005 of 042 steps.....Executing.
```

Figure 13. Creating a 240 GB Symmetrix device with EMC Solutions Enabler

Figure 13 also shows that the request for creating a RAID 5 protected device succeeds. A RAID 5 protected device on a Symmetrix VMAX storage array is striped across either four or eight physical drives. The 240 GB, RAID 5 protected devices would require approximately 80 GB of contiguous storage when it is striped across four physical drives. Since it is possible to accommodate the requirement on 146 GB drives, the configuration session shown in Figure 13 succeeds. The consumption of approximately 80 GB of storage on four of the six drives in disk group 2 is shown in Figure 14.




Figure 14. RAID 5 disk consumption for a 240 GB volume in Symmetrix Management Console

The behavior of Enginuity 5874 when creating logical devices is further elucidated in Figure 15 and Figure 16. Figure 15 shows the hard disk drives in disk group 3 on the Symmetrix VMAX storage array. Figure 16 shows the successful execution of the request to create a 240 GB, RAID 1 protected logical device on disk group 3.




Figure 15. View of a disk group containing 1 TB 7200 rpm SATA drives using SMC




Figure 16. Creating a 240 GB RAID 1 volume with EMC Solutions Enabler

The process to create the new device succeeds since the Enginuity operating environment is able to find a contiguous chunk of free space on a pair of large SATA-2 drives. This can be clearly seen in Figure 17 that shows the new logical volume fully contained within a single 1 TB SATA-2 drive.




Figure 17. Disk consumption of a 240 GB RAID 1 volume over two 1 TB SATA drives

Based on the discussion presented previously, it is clear that without proper care, it is possible to inadvertently create a new logical device on slower, large physical drives. In addition, it may be difficult or impossible to migrate large logical devices to faster drives at a later time using Virtual LUN migration technology. Therefore, the performance implication of the methodology utilized by Enginuity release version 5874 for creating logical devices should be considered before large logical devices are provisioned to VMware Virtual Infrastructure environments. Furthermore, to prevent inappropriate use of the slower disk drive technology, different hard disk drives in the Symmetrix VMAX storage array should be grouped into different disk groups, and creation of new logical devices should always be targeted to a specific disk group. It is also important to note that this discussion does not apply to Symmetrix devices such as thin devices or virtual devices, as they are solely cache devices and do not directly consume space on disks.

Auto-provisioning Groups

Introduced in Solutions Enabler version 7.0, Auto-provisioning Groups provides an easier, faster way to provision storage, reducing labor and risk of error in Symmetrix VMAX arrays running Enginuity 5874. VMware's ESX Server running on Symmetrix arrays require a highly available environment which implies multiple paths to devices. Auto-provisioning Groups was developed to make storage allocation simpler and more efficient, especially in clustered and virtualized server environments.

Mapping and masking devices in previous versions of Solutions Enabler and Symmetrix DMX arrays required a separate command for each initiator and port combination through which devices would be accessed. In Solutions Enabler 7.0, the `symaccess` command allows the user to create an arbitrary group of devices (storage group), a group of director ports (port group), and a group of host initiators (initiator group), and associate them in a masking view. In simple configurations a masking view, along with an

associated storage, port, and initiator group, can also be created for a single initiator on a single port with any number of devices using one command.

Auto-provisioning Groups includes four main parts:

- **Initiator groups** – These groups contain Fibre Channel World Wide Names (WWN) or the iSCSI IQN names of iSCSI initiators. Each initiator group can contain up to 32 different initiators. For configurations that require more than 32 initiators, initiator groups can be cascaded, that is, an initiator group can contain other initiator groups. However, an initiator group that contains references to other initiator groups cannot be cascaded. This limitation should not prove to be restrictive since the single level of cascading allows customers to create a single initiator group that has a total of 1,024 initiators.
- **Storage groups** – These groups contain Symmetrix storage device identification numbers.
- **Port groups** – These groups contain Symmetrix director port numbers.
- **Views** – This is the name of the construct that is created when an initiator group, storage group, and port group are associated with one another. A view enables a set of initiators access to a group of devices using the front-end Fibre Channel or iSCSI ports on the Symmetrix VMAX storage array.

These objects can be created using either Symmetrix Management Console or Solutions Enabler version 7.0 command line utilities.

Provisioning storage to a VMware Virtual Infrastructure

The following steps are required to provision storage from the Symmetrix VMAX system to a VMware Virtual Infrastructure environment. The process described assumes this is the first time storage is being provisioned from the Symmetrix VMAX array to the VMware Virtual Infrastructure. It also assumes that the VMware ESX Servers have been zoned to access appropriate front-end ports on the Symmetrix VMAX storage array.

1. Create devices using either Symmetrix Management Console or the Solutions Enabler command line interface.
2. Create a new storage group by executing the following command:

```
symaccess -sid <symm_id> -name <group name> -type storage create
```

For example, the command,

```
symaccess -sid 65 -name Storage_Group_Test -type storage create
```

creates the storage group, Storage_Group_Test.
3. Create a new port group. The command,

```
symaccess -sid 65 -name Port_Group_Test -type port create
```

creates the port group, Port_Group_Test.
4. An initiator group should be created next. The creation of the initiator group, Initiator_Group_Test, can be achieved by running the command,

```
symaccess -sid 65 -name Initiator_Group_Test -type init create
```
5. Once the initiator group, port group, and storage group have been created, they can be populated with the appropriate members. The process to add devices 1EA through 1F9 to storage group Storage_Group_Test is shown in Figure 18. The symaccess command to add devices to a storage group automatically detects and adds only those devices that can be addressed by the host. This can be seen in Figure 18 where only four devices (1EA, 1EF, 1F4, and 1F9) in the range 1EA through 1F9 were added to the storage group.

Figure 18 was created with a non-GA version of Solutions Enabler. The GA version of Solutions Enabler 7.0 notifies the user when metavolume members and other specialized devices such as Virtual Provisioning™ data devices are not added to the storage group.

```
[root@llicoe139 root]# symaccess -sid 65 -name Storage_Group_Test -type storage add dev 1ea:1f9
[root@llicoe139 root]# symaccess show Storage_Group_Test -type storage -sid 65

Symmetrix ID       : 000192601265
Last updated at    : 11:22:24 PM on Thu Feb 12,2009
Storage Group Name : Storage_Group_Test
Devices            : 01EA 01EF 01F4 01F9
Masking View Names
(
  None
)
```




Figure 18. Adding devices to the storage group

6. The port group, Port_Group_Test, can then be populated with the appropriate front-end ports. The procedure to add ports 7E:1 and 10E:0 to port group “Port_Group_Test” is shown in Figure 19.

```
[root@llicoe139 root]# symaccess -sid 65 -name Port_Group_Test -type port -dirport 7E:1,10E:0 add
[root@llicoe139 root]# symaccess -sid 65 show Port_Group_Test -type port

Symmetrix ID       : 000192601265
Last updated at    : 11:36:39 PM on Thu Feb 12,2009
Port Group Name    : Port_Group_Test

Director Identification
(
  FA-10E:0
  FA-7E:1
)

Masking View Names
(
  None
)
```




Figure 19. Adding ports to the port group

7. The WWPN of the Fibre Channel HBAs, or the iSCSI IQN names of the iSCSI initiators, on all of the VMware ESX Servers requiring access to storage should be added to the initiator group created in step 4. The process for populating the initiator group Initiator_Group_Test is shown in Figure 20 where the initiators have previously been stored in the file init.txt. The contents of this file are shown as an insert in Figure 20.

```
[root@licoe139 root]# symaccess -sid 65 -name Initiator_Group_Test -type initiator add -file init.txt
[root@licoe139 root]# symaccess -sid 65 show Initiator_Group_Test -type initiator
```

Symmetrix ID : 000192601265

Last updated at : 11:55:22 PM on Thu Feb 12,2009

Initiator Group Name : Initiator_Group_Test

Host Initiators

```
{
  WWN :10000000c975c6d4
  WWN :10000000c975c6d3
}
```

Masking View Names

```
{
  None
}
```

Parent Initiator Groups

```
{
  None
}
```

**Add initiators to group
(Uses text file to list initiators)**

Show Initiator Group

init.txt:

```
WWN:10000000c975c6d4
WWN:10000000c975c6d3
```

Figure 20. Adding initiators to initiator groups

8. Once an initiator group, storage group, and port group have been created and populated, they can then be associated together into a view. When a view is created, the devices included in the storage view that are not mapped to the front-end ports in the port group are automatically mapped *and* then masked to the initiators defined in the initiator group. If the devices are already mapped to the front-end ports in the port group, the devices are just masked to the appropriate initiators in the initiator group. By default, the devices are presented to the initiators in the view at the first available host address. The host addresses for the devices can be, however, changed by using the optional `-lun` parameter when creating the view. Figure 21 depicts the process for creating a view.

```
[root@licoe139 root]# symaccess create view -name Storage_View_Test -storgrp Storage_Group_Test -portgrp Port_Group_Test -initgrp Initiator_Group_Test -sid 65
[root@licoe139 root]# symaccess list view -name Storage_View_Test -sid 65 -v
```

Symmetrix ID : 000192601265

Masking View Name : Storage_View_Test

Initiator Group Name : Initiator_Group_Test

Port Group Name : Port_Group_Test

Storage Group Name : Storage_Group_Test

Show contents of View

Create View

Figure 21. Creating a view

9. After the view is created the devices are available to the host HBAs listed in the initiator group. Unlike the previous generation of Symmetrix storage array, no additional activities are required—changes go into immediate effect when the view is created.

The newly provisioned storage can be discovered in the VMware Virtual Infrastructure environment by rescanning the storage adapters of all VMware ESX Servers as shown in Figure 22. The newly discovered devices are depicted in the ESX Server context view of the EMC Storage Viewer in Figure 23.

The process described here can be further optimized by merging the step for creating initiator groups, port groups, and storage groups, with the step for adding members into the empty groups. Based on the process described here, it should be obvious to the reader that the new paradigm for provisioning storage in a Symmetrix environment simplifies the process to allocate storage for a large VMware Virtual Infrastructure environment with several VMware ESX Servers.




Figure 22. Performing a rescan of storage adapters on a VMware ESX Server




Figure 23. Newly discovered Symmetrix devices shown using EMC Storage Viewer for Virtual Infrastructure Client

Modifying existing views

After a view has been created, adding or removing objects from the view is extremely simple to do. For example, additional storage can be provisioned to an existing VMware Virtual Infrastructure environment by just adding new devices to the view's storage group. Similarly, front-end ports from the Symmetrix VMAX system can be added to an existing VMware Virtual Infrastructure environment by modifying the port group definition to include the additional ports. Furthermore, adding a new ESX Server to the VMware Virtual Infrastructure is greatly simplified with a single command that adds the new initiator to the existing initiator group. The Auto-provisioning Groups feature thus makes sharing storage among ESX Servers a much easier task.

Use Case: Expanding VMware cluster groups by sharing datastores

Frequently, a VMFS datastore that resides on a Symmetrix needs to be presented to another cluster of VMware ESX Server. In the previous generation Symmetrix system this was an arduous task that increased in complexity with the size of the new VMware ESX Server cluster group. With Auto-provisioning Groups, all that is required is for the initiators of the target ESX Servers to be added to the initiator group. As soon as that is done, the devices in the storage group associated with that view are instantly masked to the newly added initiators as long as the fabric zoning provides them access to the appropriate ports on the Symmetrix VMAX storage array.

The process is highlighted in the following example. Consider a configuration in which the datastore “Test Datastore” residing on the Symmetrix device 01EA is currently presented to the ESX Server, licoe139. The existing one-node configuration has to be expanded to meet the availability and performance requirement imposed by the business by the addition of a new VMware ESX Server, licoe134.

The first step to achieve this is to identify the WWN of the Fibre Channel HBAs in licoe134. Figure 24 shows the WWNs of the ESX host, licoe134, and the contents of the init.txt that will be used to add the initiators to the existing initiator group definition.




Figure 24. Adding initiators to a view

The WWNs of licoe134’s HBAs can be added to the initiator group, Initiator_Group_Test, by using the command:

```
symaccess -sid 65 -name Initiator_Group_Test -type init add -file init.txt
```

The resulting view, with its newly added initiators, can be seen in Figure 25.

```
[root@licoe139 root]# symaccess show Storage_View_Test view -sid 65

Symmetrix ID          : 000192601265
Masking View Name     : Storage_View_Test
Initiator Group Name  : Initiator_Group_Test

  Host Initiators
  {
    WWN : 10000000c975c6d4
    WWN : 10000000c975c6d3
    WWN : 2100001b321be03b
    WWN : 2101001b323be03b
  }

Port Group Name       : Port_Group_Test

  Director Identification
  {
    FA-10E:0
    FA-7E:1
  }

Storage Group Name    : Storage_Group_Test

Sym Dev
Name  Dir:P  Physical Device Name  Host  Lun  Attr  Cap (MB)
-----
01EA  10E:0  Not Visible                    0  (M)  163842
      07E:1  Not Visible                    0  (M)  163842
01EF  10E:0  Not Visible                    1  (M)  163842
      07E:1  Not Visible                    1  (M)  163842
01F4  10E:0  Not Visible                    2  (M)  163842
      07E:1  Not Visible                    2  (M)  163842
01F9  10E:0  Not Visible                    3  (M)  163842
      07E:1  Not Visible                    3  (M)  163842
```

Detailed contents of the view "Storage_View_Test"

Figure 25. Detailed contents of a view

The modification of the initiator group with the additional initiators automatically provides the ESX Server licoe134 access to Symmetrix device 1EA. As shown in Figure 26, the device, and the associated datastore, Test_Datastore, can be discovered on licoe134 by rescanning its storage adapters.




Figure 26. Expanding the VMware cluster by sharing a VMFS datastore

Use Case: Providing additional capacity to an existing VMware cluster group

The process to provide additional capacity to an existing VMware cluster group is similar to that described previously. The first step in the process is to create the additional devices that satisfy the storage requirement. Once the devices are created, they can be mapped and masked to the existing VMware cluster group by modifying the contents of the storage group associated to the VMware cluster group. The command,

```
symaccess -sid <symm_id> -name <group_name> -type storage add dev 100
```

adds Symmetrix device 100 from the storage array with serial number `symm_id` to the storage group `group_name`. Any initiator that is part of a view that contains the storage group `group_name` will have access to device 100.

Scalable and flexible design of Auto-provisioning Groups

The previous section discussed the new feature available in the Symmetrix VMAX storage arrays. This feature simplifies the provisioning and management of storage allocated to VMware Virtual Infrastructure. However, the simplification of the process does not come at the expense of either flexibility or scalability. The flexibility of Auto-provisioning Groups to accommodate common configurations encountered in VMware Virtual Infrastructure environments is discussed in the following sections by using two specific use cases.

Hybrid compute environments with boot from SAN for blade servers

Consider a VMware Virtual Infrastructure environment that includes a hybrid of blade and non-blade servers. The blade servers are configured without an internal hard disk necessitating a boot from the SAN environment. The proper configuration of such an environment would require that each blade server be provided with a unique Auto-provisioning Group view to its boot disk. In addition, a separate Auto-provisioning Group view that enables all VMware ESX Servers in the environment to access the shared datastores hosting the virtual machines is required.

Figure 27 shows a four-node VMware Virtual Infrastructure environment. The nodes licoe137 and licoe138 boot from the Symmetrix VMAX array. The figure shows the EMC storage details of the boot device allocated to the node licoe137 using the EMC Storage Viewer for Virtual Infrastructure Client. Further details about this tool can be found in the white paper *Using EMC Storage Viewer for Virtual Infrastructure Client* on Powerlink. It is important to note that the best practices discussed in the “Connectivity” section of this paper have been implemented when presenting the boot devices to the ESX Server.




Figure 27. Boot from SAN information using EMC Storage Viewer for Virtual Infrastructure Client

The view created for presenting the boot device to the node licoe137 is shown in Figure 28. It can be seen in this figure that the Initiator Group view, bg_e137_boot_view, used to present the boot device to the VMware ESX Server consists of the initiator group, bg_e137_hba, the port group, isv_esx_port, and storage group, bg_e137_boot_stor. The WWPN of the HBAs on the server licoe137 are included in the definition of the initiator group, bg_e137_hba. The port group, isv_esx_port, contains the FA ports 7E:0, 7E:1, 10E:0, and 10E:1. The storage group definition contains a single device, 395, that provides the boot image for the VMware ESX Server. The initiator group (not shown in the figure) consists of the host bus adapters on the server and are listed in Figure 27. A similar view, bg_e138_boot_view, was created to present the boot device for the VMware ESX Server licoe138.




Figure 28. Properties of a view through Symmetrix Management Console

Once the Auto-provisioning Group views for the boot devices have been created, the devices that are shared across the production cluster can be presented by creating a cascading initiator group that includes the initiator groups for the nodes booting from the SAN and the initiators for the nodes licoe133 and licoe134 that boot off internal disks. The process is depicted in Figure 29.

```
root@licoe069:~/host_hba
[root@licoe069 host_hba]# cat production_cluster_ig
WWN:10000000c975c05b
WWN:10000000c975c05c
WWN:2100001b321be03b
WWN:2101001b323be03b
IG:bg_e137_hba
IG:bg_e138_hba
[root@licoe069 host_hba]#
[root@licoe069 host_hba]# symaccess -sid 65 -name bg_prod_cluster_hba -type init \
> -file production_cluster_ig create
[root@licoe069 host_hba]#
[root@licoe069 host_hba]# symaccess -sid 65 show bg_prod_cluster_hba -type init | head -n
25

Symmetrix ID          : 000192601265
Last updated at       : 11:44:17 AM on Sat Mar 07,2009
Initiator Group Name  : bg_prod_cluster_hba

Host Initiators
(
  WWN :10000000c975c05b
  WWN :10000000c975c05c
  WWN :2100001b321be03b
  WWN :2101001b323be03b
  IG  :bg_e137_hba
  IG  :bg_e138_hba
)

Masking View Names
(
  None
)

Parent Initiator Groups
(
  None
)
[root@licoe069 host_hba]#
```

WWPNs of HBA not booting of SAN

Initiator Group for servers booting of SAN

Creating cascaded initiator group

Members of cascaded initiator groups are not displayed

Figure 29. Cascading initiator groups in Solutions Enabler

A new view that includes the cascaded initiator group, the appropriate port group, and a storage group consisting of the shared devices ensures that the devices shared across the production cluster are presented as needed. The process to present the shared Symmetrix devices 120 through 12F from the Symmetrix VMAX storage array 192601265 is shown in Figure 30. Figure 31 shows that the Auto-provisioning Groups software automatically assigns unique LUN addresses for the devices that do not conflict with existing address assignment used by the child initiator groups. The user can change the LUN assignment when creating the view by explicitly providing the LUN address with the optional `-lun` parameter.

Auto-provisioning Groups software, by default, automatically assigns LUN addresses for individual initiators in a cascaded initiator group configuration. This can result in the same device presented at different LUN addresses in cascaded initiator group configurations. Since VMware ESX Server version 3.5 update 1 and earlier use the LUN address when computing the signature for a VMware file system, any datastore on these devices will be viewed as a snapshot by a subset of nodes in a VMware Virtual Infrastructure that is part of a cascaded initiator group. Without changes to the advanced parameters these nodes will hide the datastores. EMC, therefore, strongly recommends the use of the optional `-lun` parameter when presenting devices from the Symmetrix VMAX storage array to VMware Virtual Infrastructure environments running VMware ESX Server version 3.5 update 1 and earlier that require the use of cascaded initiator groups.

```

root@licoe069:/tmp/se
[root@licoe069 se]# symaccess -sid 65 -name bg_prod_clus_sg -type stor \
> create dev 120:12F -reserve_id 2
[root@licoe069 se]#
[root@licoe069 se]#
[root@licoe069 se]# symaccess -sid 65 -name bg_prod_clus_view \
> -ig bg_prod_cluster_hba -pg isv_esx_port -sg bg_prod_clus_sg -reserve_id 2 create view
[root@licoe069 se]#

root@licoe069:/tmp/se
[root@licoe069 se]# symaccess -sid 65 show view bg_prod_clus_view

Symmetrix ID          : 000192601265
Last updated at       : 02:20:12 AM on Fri Mar 20,2009
Masking View Name     : bg_prod_clus_view
Storage Group Name    : bg_prod_clus_sg

Sym Dev
Name  Dir:P  Physical Device Name  Host  Lun  Attr  Cap(MB)
-----
0120  10E:0  Not Visible             0      0      32767
      07E:0  Not Visible             0      0      32767
      10E:1  Not Visible             0      0      32767
      07E:1  Not Visible             0      0      32767
0121  10E:0  Not Visible             1      1      32767
      07E:0  Not Visible             1      1      32767
      10E:1  Not Visible             1      1      32767
      07E:1  Not Visible             1      1      32767

```

Figure 30. Creating and displaying a view with cascaded initiator groups

```

root@licoe069:/tmp/se
[root@licoe069 se]# symaccess -sid 65 show view bg_prod_clus_view -ig bg_e137_hba

Symmetrix ID          : 000192601265
Last updated at       : 02:20:12 AM on Fri Mar 20,2009
Masking View Name     : bg_prod_clus_view
Initiator Group Name  : bg_e137_hba *

Sym Dev
Name  Dir:P  Physical Device Name  Host  Lun  Attr  Cap(MB)
-----
0120  10E:0  Not Visible             1      1      32767
      07E:0  Not Visible             1      1      32767
      10E:1  Not Visible             1      1      32767
      07E:1  Not Visible             1      1      32767
0121  10E:0  Not Visible             2      2      32767
      07E:0  Not Visible             2      2      32767
      10E:1  Not Visible             2      2      32767
      07E:1  Not Visible             2      2      32767

```

View for the cascaded initiator group

LUN address automatically offset for child initiator group

Figure 31. Displaying the details of a view for a cascaded initiator group

The shared devices can then be discovered by the nodes of the cluster by rescanning the HBAs. The devices visible on nodes licoe133 and licoe137 after the rescan is performed are shown in ESX Server context view of EMC Storage Viewer for Virtual Infrastructure Client in Figure 32. It can be seen that device 120 is visible on the host licoe133 with the canonical name vmhba0:4:0. The same device is accessed by the

VMware kernel on the server licoe137, at address vmhba0:1:1. This correlates to the LUN address assignments made by Auto-provisioning Groups and shown in Figure 30 and Figure 31.




Figure 32. Shared Symmetrix VMAX devices through Auto-provisioning Groups

The effect of Auto-provisioning Groups configuration described previously can be achieved by creating separate views for the two servers booting off the SAN, and another view for the rest of the servers in the environment. However, this approach adds an unnecessary level of complexity that is exacerbated as the size of the environment grows. The configuration recommended in this section requires additional steps for the initial setup but simplifies the ongoing management of a large VMware Virtual Infrastructure environment. For example, with the recommended configuration, provisioning additional shared storage to the cluster group used in this example requires manipulation of one entity – bg_prod_clus_sg. The alternative approach will require manipulation of three objects.

The use case described in this section also applies to other configurations that are similar. For example, the process described can be used for the configuration in which each VMware ESX Server needs access to a dedicated set of Symmetrix gatekeepers, and to a shared set of Symmetrix devices that host VMware file system datastores.

Isolation of workloads to a subset of the compute environment

Consider a large VMware Virtual Infrastructure environment that consists of a wide variety of server hardware. In such an environment, it is often desirable to limit critical applications to a subset of the compute environment with the flexibility of running the applications on other cluster groups that normally run less critical workloads. Such a configuration maximizes the utilization of the overall compute environment. An implementation of such an environment is shown in Figure 33. As seen in the figure, the Virtual Center Server manages the “Atlanta” data center, which contains two clusters, Production and QA/Development. The “Production” cluster, which has the two VMware ESX Servers licoa218 and licoa219, runs the production workload. The QA/Lab environment, which has the two other VMware ESX Servers licoa239 and licoa240, handles the QA and development workload.




Figure 33. Display of VMware Virtual Infrastructure supporting a varied workload

Figure 34 shows the datastores that are available in the Atlanta data center. It can be seen from the figure that the datastore, QA_Datastore, is visible to only the QA/Lab cluster. However, the datastore, Critical_Datastore, is visible to both Production and QA/Lab clusters. This configuration ensures that the QA/Lab cluster can register and run all virtual machines available on both datastores, thus allowing one to leverage the QA/Lab compute resources to run production workloads in case of an unanticipated demand. Complex VMware Virtual Infrastructure like this can be easily accommodated by Auto-provisioning Groups technology available on Symmetrix VMAX storage arrays.




Figure 34. Display showing the hosts connected to production and QA datastores

The Auto-provisioning Groups definitions to support VMware Virtual Infrastructure described previously is shown in Figure 35 and Figure 36 next. The QA and development servers that can support all of the virtual machines in the environment are defined in a separate initiator group, `bg_qa_clust_hba`. A view, `bg_qa_clust_view`, is a join of the initiator group `bg_qa_clust_hba`, port group `bg_qa_clust_port`, and storage group `bg_qa_clust_stor`. The storage group includes the device that hosts the virtual machines running the QA and development applications — that is, the device on which the datastore, `QA_Datastore`, was created.




Figure 35. The view for presenting the QA_Datastore devices to the QA cluster

In addition to the aforementioned initiator group, an additional initiator group called `bg_vi_hba` is defined. This initiator group definition contains the WWN of the HBAs on the production servers, and the initiator group, `bg_qa_clust_hba`. The storage group, `bg_vi_stor`, contains the device that hosts the datastore, `Critical_Datastore`. The view, `bg_vi_view`, created using the initiator group `bg_vi_hba`, port group `bg_qa_clust_port`, and storage group `bg_vi_stor`, ensures that the datastore, `Critical_Datastore`, is visible on all nodes in the Atlanta data center.




Figure 36. The view for presenting the production datastore to the VI environment

The creation of the Auto-provisioning Groups view using the procedure described in the previous section ensures that there are no LUN addressing conflicts. Furthermore, as the Virtual Infrastructure environment grows there is no additional complexity to manage the environment. It should be noted that, similar to the discussion in previous sections, the devices can be presented to the environment by creating non-cascading Auto-provisioning Group definitions. However, the complexity of managing such an environment increases exponentially as the VMware Virtual Infrastructure evolves.


It should be clear from the previous discussion that Auto-provisioning Groups, while simplifying storage management, also allows for greater flexibility and control. As a result, caution needs to be taken when adding or removing devices, ports, or initiators from a group. If an object is added to a group that is in multiple views, all of those views are automatically updated. This could lead to unintended consequences and cause storage to be presented to, or removed from, hosts that should not have access to them. Therefore, when modifying objects in groups associated with one or more views, it is important to fully understand the implications of the proposed changes before they are executed.

Solutions Enabler in VMware Virtual Infrastructure environments

Solutions Enabler version 6.5 and later included enhancements that enabled its use with VMware ESX Server version 3. When installed on the VMware ESX Server console, Solutions Enabler 6.5 could be used for mapping the VMkernel devices to the Symmetrix device number. This was achieved by adding optional parameters to a select set of commands that perform mapping functions. In Solutions Enabler 7.0, the support has been extended by integrating the software with the VMware Virtual Infrastructure API. This integration now enables the command `sympd` to perform mapping of the datastores hosting the virtual disks used by VMware virtual machines to Symmetrix device numbers. The functionality is similar to that offered by EMC Storage Viewer for Virtual Infrastructure Client. The procedure to use the new functionality is discussed next.

The integration with VMware Virtual Infrastructure API allows Solutions Enabler to communicate and query the VMware ESX Server or vCenter Server for information about various VMware objects. A series of steps need to be performed for Solutions Enabler to interact with the VMware API services. The first

step is to store in the SYMAPI database the credentials required for interacting with the VMware ESX Server or the vCenter Server. The command, symcfg authorization, as shown in Figure 37, should be performed in the VMware virtual machine on which the new functionality provided by Solutions Enabler is desired.



```
root@licoe084:~# symcfg authorization add -vmware -hostname 10.243.159.69 -namespace vmware/esxv2 -user root -password '
root@licoe084 ~]#
root@licoe084 ~]#
root@licoe084 ~]#
root@licoe084 ~]#
```

Figure 37. Adding credentials to the SYMAPI database for ESX Server or vCenter Server

The figure shows the process on a Linux virtual machine. The utility, vicfg.exe, located in the folder, <SE_Install_Folder>\SYMCLI\bin\shlib, where <SE_Install_Folder> is the folder in which Solutions Enabler is installed, can be used on a Windows virtual machine to perform the same function. The second step to enable Solutions Enabler to communicate with the VMware Virtual Infrastructure API is to provide the software with the public certificate associated with the VMware entity. This certificate is used by Solutions Enabler and VMware Virtual Infrastructure when exchanging information. The public certificate, rui.crt, for the vCenter server stored in the location C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\SSL should be copied to the directory <SE_Install_Folder>\SYMAPI\config as viclient_cert.pem. The <SE_Install_Folder> is the folder in which Solutions Enabler is installed. This is normally C:\Program Files\EMC\ on Windows. The certificate should be copied to the directory /var/symapi/config on UNIX servers. The public certificate rui.crt is stored in the directory /etc/vmware/ssl on the VMware ESX Server. Once these steps have been performed, the command symcfg discover automatically communicates with the VMware ESX Server or vCenter to perform the mapping. The mapping information can be displayed by executing the command symcmd list -vm, as shown in Figure 38. As seen in the figure, the virtual disk \\.PHYSICALDRIVE0 resides on Symmetrix device 287 from the array with serial number 190103232. It should be noted that the size reported in the output is not the size of the virtual disk but the size of the Symmetrix device. The figure also shows the mapping for devices that are presented as physical pass-through RDM to the virtual machine. For these devices, Solutions Enabler provides the canonical name of the device that is used by the VMware kernel when accessing the Symmetrix devices.

There is no support for mapping RDM in a virtual compatibility mode

```

C:\Documents and Settings\ganesb>syminq

```

Name	Type	Vendor	ID	Rev	Ser Num	Cap (KB)
\\.\PHYSICALDRIVE0		VMware	Virtual disk	1.0	N/A	N/A
\\.\PHYSICALDRIVE1		VMware	Virtual disk	1.0	N/A	N/A
\\.\PHYSICALDRIVE2	GK	EMC	SYMMETRIX	5773	320005C000	5760
\\.\PHYSICALDRIVE3	GK	EMC	SYMMETRIX	5773	320005D000	5760
\\.\PHYSICALDRIVE4	GK	EMC	SYMMETRIX	5773	320005E000	5760
\\.\PHYSICALDRIVE5	GK	EMC	SYMMETRIX	5773	340005B000	5760
\\.\PHYSICALDRIVE6	GK	EMC	SYMMETRIX	5773	340005C000	5760
\\.\PHYSICALDRIVE7	GK	EMC	SYMMETRIX	5773	340005D000	5760
\\.\PHYSICALDRIVE8	GK	EMC	SYMMETRIX	5773	340005E000	5760
\\.\PHYSICALDRIVE9	GK	EMC	SYMMETRIX	5874	800009B000	5760
\\.\PHYSICALDRIVE10	GK	EMC	SYMMETRIX	5874	800009C000	5760
\\.\PHYSICALDRIVE11	GK	EMC	SYMMETRIX	5874	800009D000	5760
\\.\PHYSICALDRIVE12	GK	EMC	SYMMETRIX	5874	650009A000	5760
\\.\PHYSICALDRIVE13	GK	EMC	SYMMETRIX	5874	650009B000	5760
\\.\PHYSICALDRIVE14	GK	EMC	SYMMETRIX	5874	650009C000	5760
\\.\PHYSICALDRIVE15	GK	EMC	SYMMETRIX	5874	650009D000	5760

```

C:\Documents and Settings\ganesb>sympd list -vm -sid 32
Symmetrix ID: 000190103232

```

Physical	UM	Sym	Cap (MB)
\\.\PHYSICALDRIVE0	N/A	0287	241665
\\.\PHYSICALDRIVE2	vmhba1:0:201	005C	6
\\.\PHYSICALDRIVE3	vmhba1:0:202	005D	6
\\.\PHYSICALDRIVE4	vmhba1:0:203	005E	6

```

C:\Documents and Settings\ganesb>

```

Figure 38. Using the sympd command to display virtual disk mapping information

The mapping functionality provided by Solutions Enabler can be disabled by setting the option SYMAPI_USE_VIRTUAL_BACKEND_INQUIRY to DISABLE in the file <SE_Install_Folder>\SYMAPI\config\options, where <SE_Install_Folder> is the directory in which EMC Solutions Enabler is installed on the Windows virtual machine. The equivalent file on UNIX operating systems is located in the directory /var/symapi/config. The VMAX SE product guide should be consulted for more detailed information on the new VMware Virtual Infrastructure specific functionality provided by Solutions Enabler.

Virtual LUN technology

The Symmetrix VMAX storage array enhances the Virtual LUN technology offered by DMX storage arrays. The Virtual LUN technology on the Symmetrix VMAX array can now perform transparent, nondisruptive data mobility between storage tiers for standard Symmetrix devices. This technology is based on the Enginuity Virtual RAID Architecture and requires Enginuity code level 5874.

Virtual LUN migration provides users the ability to move data between high-performance drives and high-capacity drives. This new functionality is used to either change RAID protection types, move between physical disk groups, or both. It supports the migration of regular devices or metavolumes, and can utilize configured or unconfigured disk space as the target. When the migration is performed to configured space, the target of the invocation cannot be assigned to a host. Migration is facilitated by the creation and movement of Symmetrix devices. Upon completion of the migration, the original devices are deleted and the space is freed.

Rules regarding Virtual LUNs

- Enhanced Virtual LUN Technology works with all EMC replication technologies including SRDF®, TimeFinder®, and Open Replicator.
- Migrations of devices can occur inside a single Symmetrix VMAX system.
- Once established, migrations cannot be canceled.
- Protected devices cannot be moved to unprotected devices.
- Target devices have to be unmapped and/or unmasked before they can be used for a VLUN migration session.
- Only one migration operation can be executed on any device at one time.
- Migrations are managed as a session. Once the session is established, migration is monitored and managed by the assigned session name.
- Multiple migrations can be run in parallel sessions, with a maximum of 16 sessions.
- RDF devices *cannot* be specified as the target for a migration to configured space.

Migrating VMFS datastores

Due to the ever-changing business requirements and the complex nature of a virtual environment, the SLA requirements of the underlying storage may change. This might include a need for faster response time and subsequently a need to move the VMFS volume to a faster tier of disks or a different protection scheme (RAID 5 versus RAID 1). On the other hand, changing business needs may reduce the criticality of a group of virtual machines that can be met with a lower-tier, less-expensive, and higher-capacity storage. The VLUN migration technology can meet these needs nondisruptively with no impact to the VMware Virtual Infrastructure environment. It is important to note that VMware Virtual Infrastructure provides similar functionality through the Storage VMotion software. However, this software consumes valuable CPU cycles from the VMware ESX Server to perform the operation. In addition, Storage VMotion has to be performed one virtual machine at a time and can be cumbersome for migration of a large amount of data.

Use Case: Migrating a VMFS volume from RAID 5 to RAID 1

Figure 39 shows the details of the Symmetrix device on which the VMware datastore VLUN_Test_DS was created. Figure 40 shows more detailed information about the same device including the RAID protection level and the disk group hosting the Symmetrix logical volume.

The screenshot displays the EMC Storage Viewer interface for a VMware ESX Server. The left sidebar shows a navigation tree with 'Storage' selected, containing 'Datastores', 'LUNs', and 'Targets'. The main panel is titled 'Datastores' and includes a 'Refresh' button. It contains a table with the following data:

Identification	Device	Capacity	Free	Type
licoe133:storage1	vmhba0:0:0:3	51.25 GB	50.74 GB	vmfs3
Test_Datastore2	vmhba0:1:6:1	23.75 GB	23.35 GB	vmfs3
Test_Datastore3	vmhba0:1:7:1	23.75 GB	23.35 GB	vmfs3
VLUN_Test_DS	vmhba0:2:0:1	15.75 GB	15.38 GB	vmfs3
Test_Datastore0	vmhba0:1:4:1	23.75 GB	23.35 GB	vmfs3
Test_Datastore1	vmhba0:1:5:1	23.75 GB	23.35 GB	vmfs3

Below the table is a 'Storage Details' section. It includes a table with columns: Canonical, Model, Revision, Array, Device, Type, RAID, Capacity, and # P. The data row shows: Canonical (vmhba0:2:0), Model (SYMMETRIX), Revision (5874), Array (000192601265), Device (00087), Type (RAID-5), RAID (RAID_5), Capacity (16.00 GB), and # P. (1). A red box highlights the 'Device', 'Type', and 'RAID' columns.

Figure 39. Symmetrix device before migration using EMC Storage Viewer for Virtual Infrastructure Client




Figure 40. Detailed Symmetrix device information before migration

The changes in the business now require this datastore to reside on lower-cost SATA drives with RAID 1 protection. The Symmetrix VMAX in this case has a preconfigured disk group 1 that contains low-cost 1 TB physical disk drives. To meet the business requirements a VLUN migration session can be initiated to migrate the data from RAID 5 protected 300 GB drives to disk group 1. The session can be created using either the command line interface or Symmetrix Management Console. The CLI component is utilized in this example to elucidate the process.

The migration session should first be validated using the command “symmigrate validate”. In Figure 41, the migration of device 87 from disk group 0 to disk group 1 while changing the protection level to RAID 1 is validated successfully. The file vlun.txt used in Figure 41 contains a single line to indicate that Symmetrix device 87 needs to be migrated.

```

[root@licoe139 root]# symmigrate -name ESX_VLUN -file vlun.txt -tgt_unconfig -tgt_raid1
-tgt_dsk_grp 1 validate -sid 65

Execute 'Validate' operation execution for
the device list in device file 'vlun.txt' (y/[n]) ?y

'Validate' operation execution is in progress for
the device list in device file 'vlun.txt'. Please wait...

  Validate Migration.....Started.
  Validate Migration.....Done.

'Validate' operation successfully executed for
the device list in device file 'vlun.txt'.

```

Figure 41. Validating a VLUN migration

Once the validation is successful, the migration can be initiated with the “establish” command as shown in Figure 42.

```
[root@licoe139 root]# symmmigrate -name ESX_VLUN -file vlun.txt -tgt_unconfig -tgt_raid1
-tgt_dsk_grp 1 establish -sid 65

Execute 'Establish' operation execution for
the device list in device file 'vlun.txt' (y/[n]) ?y

'Establish' operation execution is in progress for
the device list in device file 'vlun.txt'. Please wait...

  Establish Migration.....Started.
  Establish Migration.....In Progress.
  Establish Migration.....In Progress.
  Establish Migration.....In Progress.
  Establish Migration.....Done.

'Establish' operation successfully executed for
the device list in device file 'vlun.txt'.
```

Figure 42. Establishing a VLUN migration session

To query the progress of the migration, either “symmmigrate query” or “symmmigrate verify” can be used. The query command gives detailed information on the status of each migration in the session. The verify option only returns whether or not all of the migrations in the session have completed. However, the command returns a meaningful return code indicating the status of the migration and can be used in scripts. Figure 43 shows the output from the “symmmigrate query” command during the migration and after the migration is complete. Figure 44 shows an example of the output from the “symmmigrate verify” command.

Symmigrate Query during migration:

```
[root@licoe139 root]# symmigrate query -name ESX_VLUN -sid 65 -detail
Symmetrix ID: 000192601265
```

Src	Tgt	Invalid Tracks	Status SRC => TGT	Done (%)	TGT Dsk Grp	TGT Protection	Session Name
0087	N/A	219679	SyncInProg	16	01 RAID-1		ESX_VLUN

```
Total
Tracks      219679
MB(s)       13729.9
```



Symmigrate Query after migration:

```
[root@licoe139 root]# symmigrate query -name ESX_VLUN -sid 65 -detail
Symmetrix ID: 000192601265
```

Src	Tgt	Invalid Tracks	Status SRC => TGT	Done (%)	TGT Dsk Grp	TGT Protection	Session Name
0087	N/A	0	Migrated	100	01 RAID-1		ESX_VLUN

```
Total
Tracks      0
MB(s)       0.0
```




Figure 43. Querying a VLUN migration session during and after migration

Symmigrate Verify during migration

```
[root@licoe139 root]# symmigrate verify -name ESX_VLUN -sid 65
None of the session(s) with name 'ESX_VLUN' are in 'Migrated' state.
```

Symmigrate Verify after migration

```
[root@licoe139 root]# symmigrate verify -name ESX_VLUN -sid 65
All session(s) with name 'ESX_VLUN' are in 'Migrated' state.
```

Figure 44. Verifying a VLUN migration session during and after migration

Once the migration has been completed, it can be seen in Figure 45 that the protection type of the Symmetrix device has changed to RAID 1 and now resides on disk group 1, which consists of SATA drives. Also, as seen in Figure 46, the information provided by EMC Storage Viewer for Virtual Infrastructure Client has been updated to reflect the change in the protection type.




Figure 45. Detailed Symmetrix device information after migration




Figure 46. Symmetrix device information after migration using EMC Storage Viewer for VMware Infrastructure Client

Virtual Provisioning enhancements

The Symmetrix VMAX storage array provides increased storage utilization and optimization through a wide variety of new features and enhancements. One feature that has been significantly enhanced is the

Symmetrix Virtual Provisioning offering. With Enginuity release 5874, the flexibility of Virtual Provisioning and the ability to replicate thin devices have been significantly enhanced.

Thin pools can now be shrunk nondisruptively, helping reuse space to improve efficiency. When a data device is disabled, it will first move data elsewhere in the pool by draining any active extents to other enabled data devices in the thin storage pool. Once the draining is complete, that device is disabled and can then be removed from the pool.

Also with Enginuity release 5874, all tiers and RAID levels can be used and now support local and remote replication for thin volumes and pools. This includes support for RAID 1, RAID 5 (3+1), RAID 5 (7+1), RAID 6 (6+2), and RAID 6 (14+2), as well as support for TimeFinder/Clone, TimeFinder/Snap, SRDF/A, SRDF/S, SRDF/DM, Open Replicator, and Open Migrator.

The technical note *Best Practices for Nondisruptive Tiering via EMC Symmetrix Virtual LUN* provides more detailed information on the enhancements. The relevance of the enhancements in VMware Virtual Infrastructure environments will be available shortly as an update to the white paper, *Implementing Virtual Provisioning on EMC Symmetrix DMX with VMware Virtual Infrastructure*, available on Powerlink.

SRDF enhancements

Symmetrix Remote Data Facility (SRDF) is a business continuity solution that provides a host-independent, mirrored data storage solution for duplicating production data to one or more physically separated Symmetrix VMAX systems. With Enginuity 5874, new modes of operation have been added as well as upgrades to existing SRDF solutions and features. These new additions enhance the already industry-leading SRDF. The following sections list a few key enhancements to the SRDF family. For more detailed information please refer to the SRDF Product Guide available on Powerlink.

SRDF/Extended Distance Protection (SRDF/EDP)

With this new configuration, data replicates from a primary site (Site A) to a secondary site (Site B) in Synchronous mode, and then “cascades” from the secondary site to an out-of-region, extended distance, tertiary site (Site C) in the Asynchronous mode of replication. In SRDF/EDP, the secondary Site B does not contain a full data copy, instead it only retains the changed tracks received from Site A that are buffered in cache and then sent to Site C. The volumes configured in Site B are called Diskless R21 volumes (DLDEV), which have dual roles of primary (R1) and secondary (R2) volumes. These volumes have no local disk space allocated for the replicated data.

RDF Enginuity Consistency Assist (RDF-ECA)

RDF-ECA provides consistency protection for synchronous mode volumes by performing suspend operations across all SRDF/S volumes in a consistency group or a named subset of all volumes in a composite group. With Enginuity versions 5773 and earlier, in a concurrent SRDF environment, RDF-ECA could only be enabled on a single link or on both links of the concurrent SRDF relationship at the same time. Enginuity 5874 allows for the definition of an independent RDF-ECA consistency group on each of the links of a concurrent SRDF R1 volume providing the ability to disable consistency protection on one link, while still maintaining consistency on the other leg.

SRDF/A adding or removing devices

The functionality of adding and removing devices to an existing SRDF/A RDF group has been greatly enhanced in Enginuity 5874. The functionality, called the consistency exempt feature, provides the ability to dynamically add and remove volumes from an active SRDF/A session. Furthermore, it does this without affecting the state of the session or the reporting of the RDF pair state for each of the volumes in the active session that are not the target of the operation. This is achieved by marking the volumes being added or removed as “exempt” from being considered when calculating the consistency state of the volumes in the SRDF/A session, or when deciding if the SRDF/A session should be dropped to maintain dependent write consistency on the R2 side. Setting the consistency exempt flag on a volume allows the volume to be added

or removed from an active SRDF/A RDF group using either a create, delete, or move operation without requiring the other volumes in the RDF group to be suspended prior to the operation.

Remote TimeFinder/Clone to SRDF R1 restore

This feature from previous versions of the Enginuity code has been enhanced in Enginuity 5874. The feature now allows the R2 volume of a TimeFinder/Clone setup to be used to perform a SRDF restore to its partnered R1 volume while a TimeFinder/Clone restore is in progress to the R2.

SRDF/Star with the R22 device

The R22 device is a new SRDF volume type introduced in Enginuity 5874. A concurrent R2 (R22) volume is one whose two remote mirrors are paired with a different R1 volume. However, only one of the R2 mirrors may be receiving data from its corresponding R1 volume at any given time. The primary intended use for an R22 volume is to simplify failover situations and improve resiliency in SRDF/Star environments. With the introduction of the R22 volume the need to create recovery volume pairings during SRDF/Star setup is negated.

The relevance of the SRDF enhancements in VMware Virtual Infrastructure environments is discussed in the TechBook *Using EMC Symmetrix Storage in VMware Infrastructure and vSphere Environments*, available on Powerlink.

TimeFinder enhancements

TimeFinder for Symmetrix VMAX includes two modes – TimeFinder/Snap and TimeFinder/Clone. TimeFinder/Snap creates economical, pointer-based, space-saving snapshot copies of volumes, while TimeFinder/Clone creates pointer-based, full-volume clones of storage volumes for online restores, data warehouse refreshes, backups, and volume migrations. For more information please refer to the TimeFinder Product Guide available on Powerlink.

Cascaded TimeFinder/Clone

Cascaded TimeFinder/Clone, a new feature of Enginuity 5874, provides the ability to create and activate TimeFinder/Clone targets from a source device that is also a target volume of a separate TimeFinder/Clone operation. This feature allows a BCV volume, using TimeFinder/Clone emulation, to be used as a source volume for native Clone operations, while also allowing for a BCV, using Clone emulation, to be established to the target of another native Clone session.

TimeFinder/Snap Recreate

TimeFinder/Snap Recreate is a new feature introduced in Enginuity 5874. This feature simplifies the process of creating a new snap image. Prior to Enginuity 5874, taking a new point-in-time TimeFinder/Snap copy required terminating the previously activated Snap session between the source volume and the target virtual volume (VDEV). Now, following the activation of a Snap session, if a new point-in-time image is requested, the session can just be recreated and then activated again. As the recreate will lead to the replacement of the old image on the VDEV, all used tracks in the SAVE pool associated with the previous session will be released.

Use of the new TimeFinder features in VMware Virtual Infrastructure environments is discussed in the TechBook *Using EMC Symmetrix Storage in VMware Infrastructure and vSphere Environments*, available on Powerlink.

Open Replicator enhancements

Cold push from TimeFinder/Snap snapshot (VDEV)

Open Replicator cold push provides a consistent point-in-time copy of the data by making the control volume not ready to the host. When Open Replicator was first introduced in Enginuity version 5671, cold

pushes were only allowed from full volume copies such as a TimeFinder/Mirror or TimeFinder/Clone control volumes. Now, with Enginuity 5874, an Open Replicator cold push can be performed using a virtual space saving TimeFinder/Snap snapshot volume (VDEV) as the control volume. This allows users to avoid the standard copy-on-first-write penalty associated with performing a hot push from the Symmetrix, while utilizing space efficiently by not having to take a full volume copy of the data being transferred, thereby minimizing host impact. This is particularly important when users must keep applications online but don't want or need a full physical copy of the source just to perform the Open Replicator push operation.

The new feature can be of great use in VMware Virtual Infrastructure environments where a consistent copy of a business-critical application on an offsite lower tier storage array is often desired. The use of the Open Replicator enhancement for this purpose is discussed in the TechBook *Using EMC Symmetrix Storage in VMware Infrastructure and vSphere Environments*, available on Powerlink.

Conclusion

Symmetrix VMAX introduces new hardware features such as the scalable Virtual Matrix design that allows the storage array to seamlessly grow from an entry-level configuration into the world's highest-capacity storage system. A VMware Virtual Infrastructure environment backed by the power and flexibility of the Symmetrix VMAX marks a new standard for virtualized environments. Symmetrix VMAX provides improved performance and scalability for today's and tomorrow's demanding enterprise storage environments, allowing for far-reaching growth. The Symmetrix VMAX system running Enginuity Operating Environment version 5874 when used in conjunction with Solutions Enabler 7.0 provides many new features and functionality including Auto-provisioning Groups and Enhanced Virtual LUN Technology to improve data center efficiency and improve the ROI of a VMware Virtual Infrastructure environment.

References

The following documents are available on EMC's Powerlink® website:

- *Using EMC Storage Viewer for Virtual Infrastructure Client* (white paper)
- *Implementing Virtual Provisioning on EMC Symmetrix DMX with VMware Virtual Infrastructure* (white paper)
- *Using EMC Symmetrix Storage in VMware Infrastructure and vSphere Environments* (TechBook)
- *EMC Solutions Enabler Symmetrix Array Controls CLI Version 7.0 Product Guide*
- *Best Practices for Nondisruptive Tiering via EMC Symmetrix Virtual LUN Technical Note*
- *EMC Symmetrix VMAX Product Guide*

The following documentation is available on VMware.com:

- [VMware Virtual Infrastructure 3 Documentation](#)