

VMware Infrastructure 3, HP StorageWorks best practices



Executive summary.....	3
Audience	3
Information not provided.....	4
Introduction.....	4
Adaptive Enterprise	4
Virtual infrastructure	5
VMware Infrastructure 3.....	5
Using VMware Infrastructure 3 with HP StorageWorks arrays.....	5
Initial guidelines	5
Microsoft clustering	7
General SAN connectivity	7
Configuring VMFS	8
Configuring VMware VMotion.....	8
Mitigating server failure.....	8
SAN management	9
Boot from SAN	9
More information	10
Multipathing.....	10
MSA1000/1500 array.....	11
Connectivity	11
Configuration	12
Best practices.....	18
Setting a profile name.....	18
Management and monitoring	20
Array Configuration Utility (ACU).....	20
Other tools	20
EVA3000/5000 array with Virtual Controller Software (VCS) 3.028	20
Connectivity	20
More information	21

Configuration	21
Configuring failover.....	21
Viewing your current configuration	22
Tuning tips	23
More information	23
EVA4000/6000/8000 array with VCS 5.1	23
Connectivity	23
Configuration	24
Configuring failover.....	24
Viewing your current configuration	25
Tuning tips	26
Management and monitoring	26
XP128/1024 array	27
Connectivity	27
Tuning tips	27
Management and monitoring	28
More information	28
Summary	29
For more information.....	30

Executive summary

This white paper provides technical information with some of the best practices when planning or deploying VMware Infrastructure 3 with HP StorageWorks arrays.

This white paper is not designed to replace documentation supplied with individual solution components but, rather, is intended to serve as an additional resource to aid the IT professionals responsible for planning a VMware environment.

This white paper contains planning information that can help customers deploy a VMware Infrastructure 3 environment running on HP ProLiant servers, HP StorageWorks storage solutions, and HP ProLiant Essentials management software components.

Prior to reading this guide, the reader should understand the architecture of VMware ESX Server and how it virtualizes the hardware (architecture guide, VMware Infrastructure 3, How ESX Works, found at: <http://www.vmware.com/products/vi/esx/>).

All of the HP guides, white papers and technical documents for VMware Infrastructure 3 can be found at: www.hp.com/go/vmware.

To help you optimize storage in your VMware virtualized environment, HP and VMware have shared their knowledge of best practices for running VMware Infrastructure 3 with HP StorageWorks arrays, as described in this guide. Now, getting the most from your enterprise-class storage has never been easier.

IMPORTANT:

The best practices described in this guide are offered as recommendations, not requirements. The implementation of these best practices has no impact on whether or not your system may be supported by HP.

Some of these best practices may not apply to your particular environment.

Audience

This best practices information contained in this white paper is intended for solutions architects, engineers, and project managers involved with HP StorageWorks arrays with virtualization solutions. The reader should be familiar with networking in a heterogeneous environment and with virtualized infrastructures, and have a basic knowledge of VMware Infrastructure 3, and HP ProLiant servers, HP StorageWorks, HP element managers and ProLiant Essentials products.

The reader should understand the architecture of VMware Infrastructure 3 and how this product is able to virtualize hardware as outlined in the “VMware infrastructure architecture overview,” available at http://www.vmware.com/pdf/vi_architecture_wp.pdf.

Information not provided

This guide does not provide information on the following:

- Installing and configuring VMware Infrastructure 3 – for more information, refer to the VMware documents, “Installation and Upgrade Guide” (http://www.vmware.com/pdf/vi3_installation_guide.pdf) and “Server Configuration Guide” (http://www.vmware.com/pdf/vi3_server_config.pdf)
- Installing and configuring HP storage administration software



Storage System Scripting Utility (SSSU) is not supported on a VMware ESX Server 3.0 host.

-
- Installing and configuring HP servers and HP StorageWorks arrays – for in-depth information, refer to the “HP StorageWorks SAN design reference guide” at <http://h20000.www2.hp.com/bc/docs/Support/SupportManual/c00403562/c00403562.pdf>.



This guide does not discuss Internet SCSI (iSCSI) or Network Attached Storage (NAS), which have not, at this time, been certified by VMware and HP for use with VMware Infrastructure 3.

For more information on iSCSI and NAS, refer to the “SAN Configuration Guide” at http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf.

Introduction

A key component of an Adaptive Enterprise, as envisioned by HP, is the virtualization of resources. This section provides more information on these concepts.

Adaptive Enterprise

The HP Adaptive Enterprise strategy combines industry-leading solutions, services, and products from HP and partners that can help organizations quickly turn challenges into opportunities. This strategy is based on four design principles – simplification, standardization, integration, and modularity – which, when applied consistently across business processes, applications, and infrastructure, will ultimately lead to an organization that can adapt to – even embrace – change. These design principles are applied to individual elements of the IT infrastructure and the entire infrastructure itself; in this way, organizations can create consistent building blocks that can be combined as needed.

An Adaptive Enterprise is not a single product; it cannot be purchased “off the shelf”. It is a philosophy designed to make an organization agile and easily adaptive to changing business needs.

Virtual infrastructure

Virtualization is one of the cornerstones to an Adaptive Enterprise. The primary benefit to virtualization may indeed be consolidation; however, a virtualized infrastructure can be beneficial in many other ways. For example, because an entire operating environment can be encapsulated in several files, that environment becomes easier to control, copy, distribute, and so on. If an organization virtualizes an operating system, its applications, configuration settings, and other desirable elements, that entire operating environment – known as a Virtual Machine (VM) – can be rolled out anywhere in the organization to maintain business continuity. To maximize availability, emerging technologies can allow VMs to automatically migrate from a potentially failing host to another virtualized platform – with little or no user intervention.

VMware Infrastructure 3

VMware Infrastructure 3 is virtual infrastructure software for portioning, consolidating, and managing computing resources in mission-critical environments. VMware Infrastructure 3 provides a highly scalable Virtual Machine (VM) platform with advanced resource management capabilities.

Using VMware Infrastructure 3 with HP StorageWorks arrays

HP has tested and certified a broad range of HP StorageWorks storage arrays for use with VMware Infrastructure 3. This guide provides best practices for the following arrays:

- Modular Smart Array (MSA) 1000/1500
- Enterprise Virtual Array (EVA) 3000/5000
- EVA4000/6000/8000
- XP Disk Array 128/1024

Initial guidelines

HP offers the following guidelines for activities to be performed before deploying VMware Infrastructure 3 on to your SAN or HP storage:

- Plan the SAN or disk architecture. The planning phase is critical; many organizations have discovered after they created their LUNs that they had more storage capacity than needed, driving up deployment costs.
If you are unaware of your storage needs, please consult an HP storage specialist. Alternatively, you can use the HP ProLiant server sizer for VMware (www.hp.com/solutions/vmware) or HP ProLiant Essentials Performance Management Pack (<http://h18004.www1.hp.com/products/servers/proliantessentials/valuepack/PMP/downloads.html>).
- If you wish to optimize your VMs, storage plays an important role. HP offers a range of guidelines:
 - **High-performance/high-availability**
You should provide built-in snapshots to facilitate backups and point-in-time restores, replication (such as HP StorageWorks Continuous Access), service processor redundancy, and fibre drives.
 - **Mid-performance/lower-availability**
You should provide limited snapshots, limited service processor redundancy, and SCSI drives.
 - **Low performance**
There is no need to provide snapshots or clones, or configure service processor redundancy. Low-end SCSI or SATA drives are sufficient.

- Server failover designs with SANs are important. All LUNs used by cluster VMs must be visible to all ESX server hosts.
For more information, refer to the “VMware infrastructure architecture overview” at http://www.vmware.com/pdf/vi_architecture_wp.pdf.
- Verify all hardware components in the storage solution are running the latest supported firmware for VMware Infrastructure 3. For more information, contact HP technical support at www.hp.com/support.
- Verify that your HP StorageWorks array(s) have been tested and certified with VMware Infrastructure 3 by reviewing the VMware Hardware Compatibility List (HCL) document, “HCL: Storage / SAN Compatibility Guide for ESX Server 3.0,” available at http://www.vmware.com/pdf/vi3_san_guide.pdf, or by contacting HP technical support.
- To save time when adding multiple HBAs, deploy and configure all HBAs at the same time. If you add HBAs sequentially, you must reboot your VMware Infrastructure 3 host each time.
- When initially setting up an ESX server with VMware Infrastructure 3, disconnect your external disks so that they will not be visible to the server and cannot be overwritten accidentally.
- If you are implementing a diagnostic partition with HP BladeSystem servers, consider the following guidelines:
 - HP ProLiant server blades that boot from SAN can share a single diagnostic partition on a SAN LUN.
 - If multiple ESX servers are using a single diagnostic partition as a LUN, then zone the LUN so that all the hosts can gain access to it.
 - Each server requires at least 100 MB and must be mapped to a diagnostic slot.
 - 16 diagnostic slots are recommended.
 For more information on diagnostic partitions, refer to the VMware document, “SAN Configuration Guide,” at http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf.
- Table 1 provides guidelines for datastore partitions.

Table 1: Datastore partitions

Partition name	File system format	Size	Description
/boot	ext3	100MB	The boot partitions stores files required to boot ESX Server.
/	ext3	2560MB	Called the “root” partition, this contains the ESX Server operating system and Web Center files. Allocate an additional 512MB if you plan to use this server for scripted installations.
NA	Swap	544MB	The swap partition is used by the service console and tools like the HP Insight Management agents.
vmkcore	vmkcore	100MB	This partition serves as a repository for the VMkernel core dump files in the event of a VMkernel core dump.
VMFS	VMFS-3	1200MB+	The VMFS file system for the storage of virtual machine disk files. Must be large enough to hold your VM disks.
/home (optional)	ext3	512MB	Storage for individual users.
/tmp (optional)	ext3	1024MB	Partition used for temporary storage.
/var (optional)	ext3	1024MB	Partition is used for log file storage. HP recommends creating a /var partition to prevent unchecked log file growth from creating service interruptions.

As shown in Table 1, an ESX server's local boot volume requires three specific partitions. In addition, a local or remote VMFS3 partition is required to store your VMs and a VMKCORE partition is required to provide core dumps for technical support.

For information on optional partitions, refer to the VMware document, "Installation and Upgrade Guide," at http://www.vmware.com/pdf/vi3_installation_guide.pdf.

Microsoft clustering

Clustering VMs between ESX server hosts requires disks to be shared. To configure a QLogic HBA for clustering support, enter the Fast!UTIL utility during server POST, then select the desired HBA. Select **Configuration Settings**→**Advance Adapter Settings** and ensure the following settings are configured:

Enable LIP reset	No
Enable LIP full login	Yes
Enable target reset	Yes

For more information, refer to the VMware document, "Setup for Microsoft Cluster Service," at http://www.vmware.com/pdf/vi3_vm_and_msocs.pdf.

General SAN connectivity

VMware Infrastructure 3 and HP StorageWorks storage combine to provide a robust solution for consolidation, reliability, and disaster recovery. This section provides general guidelines for connecting VMware Infrastructure 3 to an HP StorageWorks SAN.

- Ensure each LUN has the appropriate RAID level and storage characteristics for the VMs that will use it.
- Carefully determine the number of VMs you can place on a single LUN in your particular environment. Note that, with new VMware Infrastructure 3 capabilities such as iSCSI and NAS, this determination has become even more complex.
While VMware Infrastructure 3 supports up to 100 VMs per LUN, 32 VMs per LUN may often be acceptable. However, HP recommends being more conservative and using 8 – 10 VMs per LUN. In a NAS or iSCSI environment, HP recommends using 8 – 16 VMs per LUN.
If your VMs are not heavily impacting I/O, you may be able to exceed the conservative recommendations above.
- To optimize array performance, spread I/O loads over the available paths to the storage.
- HP suggests configuring one Virtual Mirror File System (VMFS) volume per LUN. With VMFS-3 you no longer have to set accessibility.
- In an environment with HP BladeSystem servers, a shared diagnostic partition is appropriate.
- Local SCSI controllers and Fibre Channel HBAs can be allocated to the VMkernel. However, it may be preferable to use local RAID solutions for the Service Console and Fibre Channel HBAs for external storage such as SANs.
- Raw Device Mapping (RDM) is recommended for accessing any RAW disk for VMware Infrastructure 3. For more information, refer to the VMware documents, "Using Raw Device Mapping" (http://www.vmware.com/pdf/esx25_rawdevicemapping.pdf) and "Server Configuration Guide" (http://www.vmware.com/pdf/vi3_server_config.pdf).

- Use multipathing, which provides continued access to SAN LUNs in the event of hardware failure. HP recommends that each LUN should be presented to the same LUN number to all ESX server hosts.
- Ensure that the appropriate BusLogic or LSILogic driver is selected and that the driver's queue depth setting is large enough. Check the "Guest Operating System Guide" for details at http://www.vmware.com/pdf/GuestOS_guide.pdf
- By default, VMware Infrastructure 3 uses the QLogic **QDepth** default value of 32 to act as an execution throttle.

As a best practice, **QDepth** should be limited to around eight to retain some headroom for expansion. Refer to VMware Knowledge Base article 1267 on how to adjust the QDepth for performance tuning.

In general, the following formula applies to each path between the host machine and the array:

$$[\text{Total number of LUNs}] * \text{QDepth} < [\text{Array queue depth}]$$

Refer to VMware Knowledge Base article #1267 at

www.vmware.com/support/kb/enduser/std_adp.php?p_faqid=1267 for more information on tuning **QDepth**.

- The default queue depth value of 32 is satisfactory for Emulex HBA cards.
- If you are using Microsoft® Windows® 2000 or Windows Server 2003 check your operating system timeout values. For a VMware environment, the **DiskTimeOutValue** must be set to **60** seconds.
- Dynamic load balancing is currently not supported by VMware Infrastructure 3. However, for active/active arrays with **fixed** failover policy, load balancing can be achieved by manually assigning a LUN to a specific path. For more information, refer to the VMware document, "Server Configuration Guide," at http://www.vmware.com/pdf/vi3_server_config.pdf.

Configuring VMFS

- A LUN should contain only a single VMFS volume.
- Maintain test VMs on separate VMFS volumes from production VMs.
- Maintain VirtualCenter templates on separate VMFS volumes from virtual disk files.
- Allocate a large LUN for the use of multiple VMs; set up this LUN as a VMFS.
- To take advantage of hardware snapshot functionality in the disk array, use Raw Device Mapping (RDM) for VMs.

Configuring VMware VMotion

- LUNs holding VM virtual disks must be visible to both source and destination servers. This is very important when you plan for Distributed Resource Services (DRS) and High Availability (HA).

Mitigating server failure

- Ensure all host servers can access the same shared storage.
- All LUNs used by clustered VMs must be visible to all host servers. Remember that, even though a LUN is accessible to a particular host, all the VMs on that host do not necessarily have access to all data on that LUN.

SAN management

A range of tools is available for VMware Infrastructure 3 and SAN administrators wishing to administer or monitor the SAN. These tools include:

- **Command View**

HP recommends using HP StorageWorks Command View EVA or XP to set up, configure, and manage your array. The latest version of Command View is supported on Windows servers. Please review the system and OS requirements for Command View.

- <http://h18006.www1.hp.com/products/storage/software/cmdvieweva/index.html>
- http://www.hp.com/products1/storage/products/disk_arrays/xpstoragesw/commandview/index.html



You should review the system and OS requirements for Command View.

- **ACU**

For MSA-family arrays, HP recommends the use of the Array Controller Utility (ACU) for setup, configuration, and management.

This utility provides a command line interface.

- **HP Storage Essentials**

Storage Essentials can provide an in-depth analysis of your HP storage usage. For more information, visit <http://h18006.www1.hp.com/products/storage/software/e-suite/index.html>.

- **HP SAN Visibility Tool**

SAN Visibility is a free tool designed to help you better design, administer, and manage your SAN. Combining technical depth with expert advice on how to optimize new and existing SANs, this tool must be run from a Windows server or supported OS. For more information, visit

<http://h18006.www1.hp.com/storage/saninfrastructure/index.html>.

Boot from SAN

Traditionally, HP ProLiant servers boot operating systems from internal SCSI and ATA storage devices. Optionally, however, you can take advantage of external boot support offered by particular HBAs and arrays to eliminate the need for these server-based internal boot devices.

Booting from an external device provides multiple benefits:

- Improved disaster recovery
- Reduced backup window through the SAN
- Additional SAN-managed features

VMware Infrastructure 3 has the capability to boot directly from a SAN device with either RDM or a VMFS-3 file system. For more information, refer to the VMware document, "SAN Configuration Guide" (http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf) and VMware knowledge base articles 1458, 1463, and 1540 (http://www.vmware.com/support/kb/enduser/std_alp.php).

When implementing boot from SAN, follow these guidelines:



Failure to comply with these guidelines may result in data loss.

- The physical host and HBA(s) must be compatible with VMware Infrastructure 3. Refer to the VMware HCL, "HCL: Storage / SAN Compatibility Guide for ESX Server 3.0," available at http://www.vmware.com/pdf/vi3_san_guide.pdf, for more information on SAN devices certified by VMware.
- The HBA associated with the boot-from-SAN LUN must be shared between the Service Console and the VMkernel.
- Local disk drives must be removed from the physical host. No local storage may be present.
- Boot from SAN capability must be deployed using the installation CD or a scripted install. There is currently no support in HP ProLiant Essentials Rapid Deployment Pack (RDP) for deploying boot from SAN.
For more information, refer to the VMware document, "Installation and Upgrade Guide," at http://www.vmware.com/pdf/vi3_installation_guide.pdf.
- The installation LUN must be the lowest-numbered of the LUNs available to the physical host. In addition, the lowest-numbered path to the boot LUN must be the active path.
- The HBA serving the boot-from-SAN LUN must be the lowest-numbered HBA with available LUNs.
- The storage processor presenting the LUN must be the lowest-numbered storage processor accessible to ESX Server.
- The HBA BIOS must be enabled; the HBA must be configured to boot from SAN. Refer to your hardware documentation for more information on configuring the HBA.
- A boot from SAN environment does not support Direct Connect or FCAL. Switched fabric must be used.
- Zone LUNs so that a particular ESX Server instance is unable to access boot LUNs associated with other VMware Infrastructure 3 instances.
- During the installation, zone LUNs so that an ESX Server instance can only access its boot LUN. Later, you can carefully provide access to data LUNs but only after ensuring that they do not have lower numbers than the boot LUN.

More information

Refer to the VMware document, "SAN Configuration Guide" at http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf and VMware knowledge base articles [1458](#), [1463](#) and [1540](#) for more information on implementing boot from SAN.

Multipathing

By default, VMware Infrastructure 3 only uses one path from the host to a given LUN at any given time. If the path were to fail, ESX Server would detect and select another path for failover.

For additional information, please review the SAN compatibility guide at <http://www.vmware.com/vmtn/resources/cat/91,100>.

Use the following settings:

- **MSA1000/1500** – **Active/Passive** set to **MRU**
- **EVA3000/5000** – **Active/Passive** set to **MRU**

- **EVA3000/5000 – Active/Active** set to **Fixed**
- **EVA4000/6000/8000 – Active/Active** set to **Fixed**



Guest multipathing software cannot be used to perform I/O load balancing to a single LUN.

In the next few sections, HP isolates particular HP StorageWorks array products and provides best practices for each.

For more information on supported products, refer to the VMware HCL document, “HCL: Storage / SAN Compatibility Guide for ESX Server 3.0,” available at http://www.vmware.com/pdf/vi3_san_guide.pdf.

MSA1000/1500 array

This section provides basic guidelines for connecting an HP StorageWorks MSA array to an ESX server, offers configuration best practices, and identifies utilities and tools that can help you better manage the array.

Connectivity



Array performance can be significantly impacted by the number of users and the type of application. Please refer to HP SAN Configuration guide for assistance.

- Ensure that the array is running the latest firmware:
 - MSA1000: 4.48 minimum
 - MSA1500: 5.02 minimum
- Once the array has been configured and Selective Storage Presentation (SSP) functionality has been committed, scan the SAN to identify currently-presented LUNs. Make sure the Service Console is aware of the LUN changes.
- HP recommends configuring all physical hosts connected to the array with two HBAs and use native VMware multipathing (MPIO).
For example, consider a scenario where one physical host connects to the array via a single path; additional hosts connect via dual paths. Should the array’s primary controller become inactive, the single-path host would lose access to the array while the dual-path hosts would not.
- The MSA array can support up to 20 server nodes – either as single or clustered nodes. With 20 nodes connected, there can be a maximum of three clusters.
- HP recommends not starting your hosts or VMs at the same time as this may create I/O contention that would impact the performance of the array and hosts. To prevent this, configure your VMs to automatically delay start ups. For more information, please refer to the VMware document, “SAN Configuration Guide” at http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf.

Configuration

The MSA1500 array supports multiple options for disk layout. The following figures represent HP recommendations for configuring SCSI arrays (using the MSA30 enclosure) and SATA arrays (using the MSA20 enclosure).



Currently, you cannot mix MSA30 and MSA20 enclosures on the same MSA1500cs I/O module. A single MSA30 enclosure or two MSA20 enclosures can be attached per SCSI I/O module. The MSA1000 array only supports SCSI drives.

Figure 1 shows a MSA20 enclosure with 12 drives. The drives in this array are configured optimally with three LUNs, each shown in a different color. A single I/O module is connected to this enclosure, creating a single bus.

Figure 1: SATA array¹



¹ The MSA20 enclosure only supports a single bus and does not support SAS drives.

Figure 2 shows four MSA20 enclosures, each with three SATA drives. Each enclosure is connected to a different I/O module, creating four buses. The optimal configuration is shown, with LUNs spread equally across each bus.

Figure 2: SATA array with four buses – optimized




Figure 3 shows a larger SATA array, again with four buses. Each bus is configured with two enclosures.

Figure 3: Larger SATA array with four buses




Figure 4 shows the same basic SATA array as Figure 3, but optimally configured. In this example, LUNs are spread equally across each bus.

Figure 4: Larger SATA array with four buses – optimized




Figure 5 shows a mixed array, consisting of MSA20 (SATA) enclosures and MSA30 (SCSI) enclosures.

Figure 5: Mixed SATA and SCSI array with four buses





Figure 6 shows the same mix of enclosures as Figure 5, but optimally configured with LUNs spread equally across each bus.

Figure 6: Mixed SCSI and SATA array with four buses – optimized



Best practices

- For up-to-date information on supported configurations, refer to the VMware HCL document, "HCL: Storage / SAN Compatibility Guide for ESX Server 3.0," available at http://www.vmware.com/pdf/vi3_san_guide.pdf.
- You can have up to four I/O modules, each with two ports. HP recommends assigning a dedicated I/O module to each enclosure.
If you have four MSA20 enclosures, you should ideally configure a dedicated I/O module to each. If you were to add a fifth MSA20 enclosure, you would have to use the second port on one of the I/O modules.
Always populate Port A first. When all available Port A bays have been populated, populate Port B bays, starting with Bus 0
- A host server can support either single or redundant paths, but not both.
- When configuring **Disk.SchedNumReqOutstanding**, verify that the appropriate queue depth for your storage adapter is set. For more information on configuring this value for QLogic adapters, refer to VMware knowledge base article #1267 at www.vmware.com/support/kb/enduser/std_adp.php?p_faqid=1267.



Remember to reboot the server.

Setting a profile name

The following guidelines for setting a profile name are based on information contained in the VMware Infrastructure 3 release notes, which suggest using the MSA1000 command line interface to create a static connection to the array.

Refer to HP StorageWorks MSA1000 documentation (available at <http://h18006.www1.hp.com/products/storageworks/msa1000/documentation.html>) to learn more on how to install and configure the command line interface.



Connection settings cannot be created using HP Array Configuration Utility.

Perform the following steps:

1. Connect the MSA1000 command line interface to the MSA1000.
2. Verify that the Fibre Channel network between the MSA1000 and the host machine is functioning.

3. Start the command line interface. At the prompt, enter:

SHOW CONNECTIONS

The output displays a connection specification for each Fibre Channel WWNN/WWPN attached to the MSA1000, for example:

Connection Name: <unknown>
Host WWNN = 22222222-22222222
Host WWPN = 33333333-33333333
Profile Name = Default
Unit Offset 0
Controller 1 Port 1 Status = Online
Controller 2 Port 1 Status = Online

Ensure that the host's WWNN and WWPN show the correct connection for each HBA.

4. Create a static connection:

ADD CONNECTION ESX_CONN_1 WWNN=22222222-22222222 WWPN=33333333-33333333 PROFILE=LINUX

5. Verify the connection:

SHOW CONNECTIONS

The output displays a single connection with the 22222222-22222222/33333333-33333333 WWNN/WWPN pair and its **Profile Name** set to **Linux**:

Connection Name: ESX_CONN_1
Host WWNN = 22222222-22222222
Host WWPN = 33333333-33333333
Profile Name = Linux
Unit Offset = 0
Controller 1 Port 1 Status = Online
Controller 2 Port 1 Status = Online

IMPORTANT:

Ensure that WWNN = 22222222-22222222 and WWPN = 33333333-33333333 display a single connection. There should not be a connection set to "Unknown" for WWNN = 22222222-22222222 and WWPN = 33333333-33333333.

6. Add static connections with different values for **Connection Name** for each WWNN/WWPN on the host machine.

Management and monitoring

Array Configuration Utility (ACU)

With ACU, you can have your array controller set up and ready to use in minutes by answering a few basic questions in the configuration wizards. You can easily change your array configurations as your storage needs grow with innovative features such as Online Capacity Expansion, Logical Drive Capacity Extension, and RAID Level Migration. You can also fine tune the controller's performance by personalizing your read/write cache ratio, stripe size, and expand/rebuild priority levels.



ACU GUI does not work in a VMware host or guest operating system. You must install ACU on a Windows client with a supported HBA.

HP recommends using the latest version of ACU. Review the release notes to learn more about the fixes and enhancements made to ACU and to firmware for the MSA array.

Other tools

Other tools are also available to help you monitor and manage your array; selecting a particular tool is a matter of personal choice. These tools include the following:

- VMware **esxtop** utility can help IT staff understand the state of the host machine in a virtual infrastructure. By default, this utility's display is updated every five seconds; however, the update frequency can be changed using the **d** or **s** command.
VMware Infrastructure 3 system resource variables (such as **Disk Statistics**) can be used to track physical disk usage, including total reads or writes per second and MB read and written per second. A display of physical disk information is provided per LUN per target per HBA. The display can be toggled using the **d** command.
For more information, refer to the VMware document, "Resource Management Guide" at http://www.vmware.com/pdf/vi3_esx_resource_mgmt.pdf.
- HP SIM together with Storage Essentials and Performance Management Pack can provide a more in-depth view of the array.
- Use real-time data collected from the host server's **/proc** nodes, the web-based **vmkusage** utility can display historical graphs that show physical server, VMware Infrastructure 3, and VM statistics. By default, these graphs show most recent, daily, and weekly data.

EVA3000/5000 array with Virtual Controller Software (VCS) 3.028

Connectivity

Although HP StorageWorks EVA arrays are designed to support a wide range of applications, some configuration options may influence performance, usable capacity, or availability. As a result, you should understand the following options so you can optimize your array for a specific application:

- Number of disk groups
- Type of disk group
- Type and number of disks in a disk group
- Vraid levels (0, 1, or 5)

- Disk failure protection level (none, single, or double)
- Cache settings

It may not be possible to simultaneously optimize a particular configuration for cost, performance, and availability; indeed, there may be conflicting recommendations. For example, Vraid0 is clearly the best choice from a cost perspective because nearly all storage is available for user data; however, Vraid0 offers no protection from a disk failure so either Vraid1 or Vraid5 would be a better choice from an availability perspective. By comparison, other tradeoffs are minor but are sometimes necessary to best meet the needs of the desired configuration.



With VCS 3.028 you still need to configure a custom-mode number setting or 16-character hexadecimal value.

More information

For up-to-date information on supported configurations, refer to the VMware HCL document, "HCL: Storage / SAN Compatibility Guide for ESX Server 3.0," available at http://www.vmware.com/pdf/vi3_san_guide.pdf.

Configuration

- As with other HP StorageWorks arrays, a best practice before beginning to configure and implement an EVA3000/5000 array is to verify that the latest firmware revisions are installed. In this case, ensure that the latest supported version of HP StorageWorks Virtual Controller Software (VCS) is always applied; HP recommends VCS 3.0.28.
- In Command View EVA, set the host OS name parameter to **Custom**.
- For an active/passive EVA3000/5000 array, use host mode type **000000002200282E**.

Configuring failover

To configure a highly-available SAN failover configuration (such as that shown in Figure 7) that includes supported EVA models equipped with two Storage Processors (SPs), each host machine requires the following:

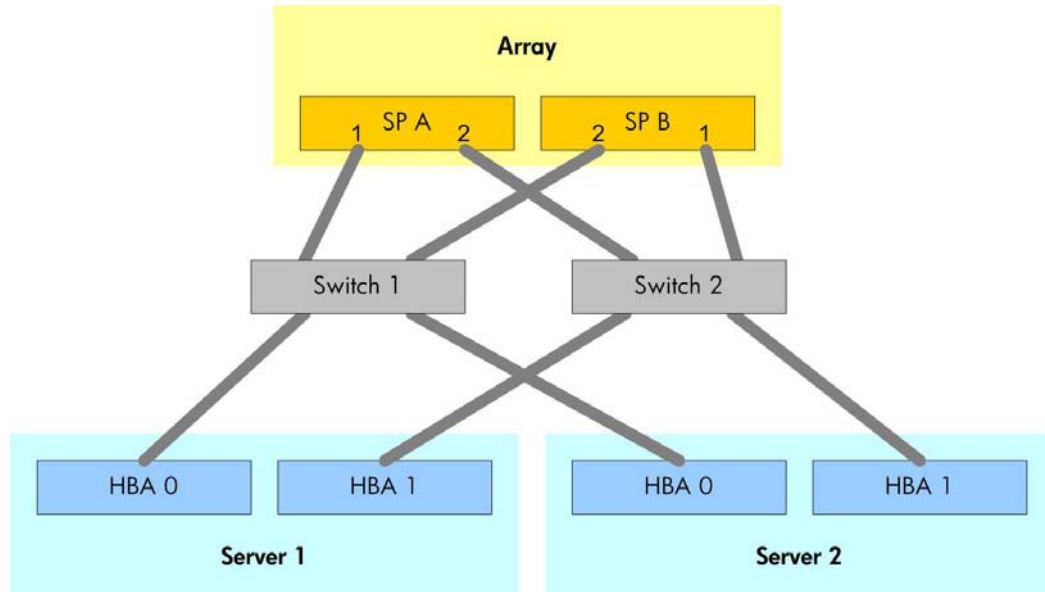
- Two Fibre Channel HBAs (such as devices from QLogic or Emulex)
- Two fibre switches
- Two SPs, each with at least two ports connected to the SAN

Single point of attachment VMware Infrastructure 3 servers must be configured (zoned) so they attach to the EVA through a common controller and host port in order to avoid controller "thrashing". Thrashing is a condition where both controllers alternate as the active or owning controller of a shared virtual disk. This behavior can be observed using **EVAPERF** to monitor shared LUNs.

Frequent thrashing can increase controller CPU utilization, which, in turn, impacts lower-priority backend processes such as leveling or reconstruct.

For more information on thrashing and on configuration in general, refer to the VMware document, "SAN Configuration Guide" at http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf.

Figure 7: Typical failover configuration



Use the following connection settings:

- HBA 0 on each server connects to Switch 1.
- HBA 1 on each server connects to Switch 2.
- Switch 1 connects Port 1 on SP A and Port 2 on SP B.
- Switch 2 connects Port 2 on SP A and Port 1 on SP B.

Connect the SP A ports to lower switch port numbers than those used to connect to the SP B ports. This allows SP A ports to be part of the first paths seen by the HBAs.

This configuration provides two paths from each HBA. LUNs must be owned by the SP to which the preferred and active paths are connected – SP A in Figure 7.

Note:

The configuration shown in Figure 7 assumes that the switches are not connected via Inter-Switch Link (ISL) in one fabric.

Viewing your current configuration

To view your current multipathing configuration, log on to your host machine as **root** and enter:

```
esxcfg-mpath -l
```

The output lists all known paths to each SAN LUN. An asterisk (*) indicates that the path is the current active path; # specifies the preferred path from the host machine to the LUN.

Tuning tips

- In an EVA array, HP recommends laying out your disks vertically. Distribute the disks across all enclosures, allowing the array to organize the Redundant Storage Sets (RSS) groups.
- In an environment with a mixture of SCSI and SATA drives, HP recommends storing your critical data on a dedicated RAID 1 SCSI disk group. You should configure SATA drives on their own disk group to be used for non-critical, less-accessed data.
- In a performance environment, HP recommends configuring one large disk group that can take advantage of all the disks in the array. If you need more than one disk group, HP recommends a maximum of three disk groups to optimize LUN manageability.
- HP recommends running HP StorageWorks Command View software on your host server to monitor the array.
- HP suggests consulting with your HP storage specialist or HP Services representative to help you identify tools that can enhance your experience with HP StorageWorks arrays and VMware Infrastructure 3.

More information

- For more information on resolving general performance issues, refer to the VMware document, “Basic System Administration,” at http://www.vmware.com/pdf/vi3_admin_guide.pdf.
- For more information on optimizing storage performance, refer to the VMware document, “SAN Configuration Guide,” at http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf.

EVA4000/6000/8000 array with VCS 5.1

Connectivity

For information on connecting EVA4000, EVA6000, or EVA8000 arrays to VMware Infrastructure 3, refer to the HP installation and reference guide, “HP StorageWorks 4000/6000/8000 EVA connectivity for VMware” at <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00509268/c00509268.pdf>.



With VCS 5.1, you no longer need to configure a custom-mode number setting for the OS type with HSV 200 and 210 controllers.

Although HP StorageWorks EVA arrays are designed to support a wide range of applications, some configuration options may influence performance, usable capacity, or availability. As a result, you should understand the following options so you can optimize your array for a specific application:

- Number of disk groups – as a best practice, HP recommends not configuring more than three disk groups
- Type of disk group
- Type and number of disks in a disk group – using all available online disks delivers performance gains but there is a cost trade-off
- Vraid levels (0, 1, or 5) – for performance, HP recommends RAID 1
- Disk failure protection level (none, single, or double)
- Cache settings

It may not be possible to simultaneously optimize a particular configuration for cost, performance, and availability; indeed, there may be conflicting recommendations. For example, Vraid0 is clearly the best choice from a cost perspective because nearly all storage is available for user data; however, Vraid0 offers no protection from a disk failure so either Vraid1 or Vraid5 would be a better choice from an availability perspective. By comparison, other tradeoffs are minor but are sometimes necessary to best meet the needs of the desired configuration.

For up-to-date information on supported configurations, refer to the VMware HCL document, "HCL: Storage / SAN Compatibility Guide for ESX Server 3.0," at http://www.vmware.com/pdf/vi3_san_guide.pdf.

Configuration

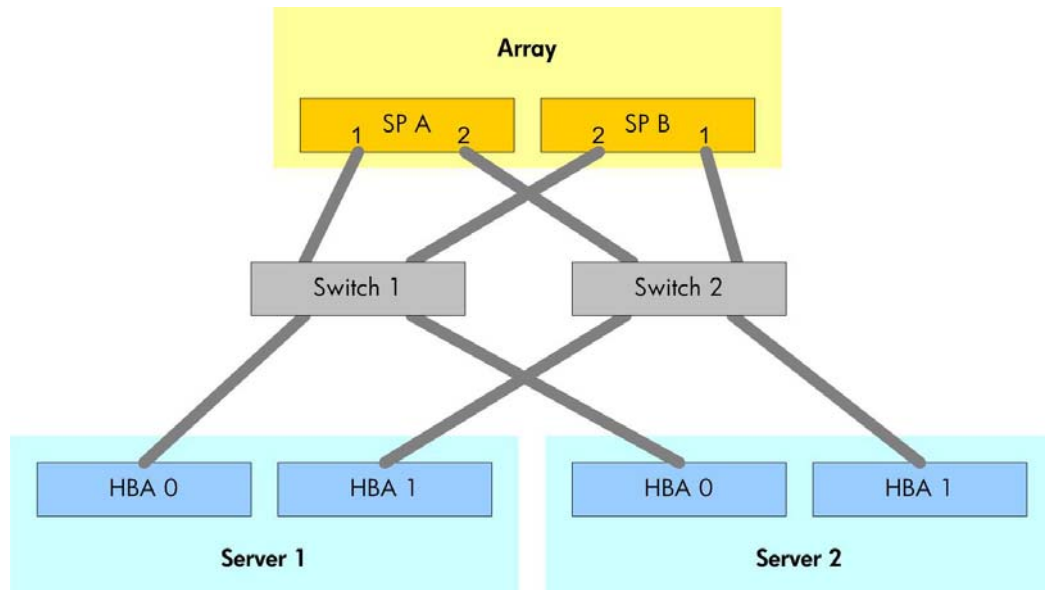
- As with other HP StorageWorks arrays, a best practice before beginning to configure and implement an EVA4000/6000/8000 array is to verify that the latest firmware revisions are installed. In this case, ensure that the latest supported version of HP StorageWorks Virtual Controller Software (VCS) is always applied; HP recommends VCS 5.1.
- In Command View EVA, set the host OS name parameter to **VMware**.
- For EVA4000/6000/8000 active/active arrays with firmware 5.031 or below, use host mode type **00000000220008BC**.
- Since EVA4000/6000/8000 controllers are active/active, HP recommends setting the Multiple Path I/O (MPIO) feature to **Fixed**.

Configuring failover

To configure a highly-available SAN failover configuration (such as that shown in Figure 8) that includes supported EVA models equipped with two Storage Processors (SPs), each host machine requires the following:

- Two Fibre Channel HBAs (such as models from QLogic or Emulex)
- Two fibre switches
- Two SPs, each with at least two ports connected to the SAN

Figure 8: Typical failover configuration



Use the following connection settings:

- HBA 0 on each server connects to Switch 1.
- HBA 1 on each server connects to Switch 2.
- Switch 1 connects Port 1 on SP A and Port 2 on SP B.
- Switch 2 connects Port 2 on SP A and Port 1 on SP B.

Connect the SP A ports to lower switch port numbers than those used to connect to the SP B ports. This allows SP A ports to be part of the first paths seen by the HBAs.

This configuration provides two paths from each HBA. LUNs must be owned by the SP to which the preferred and active paths are connected – SP A in Figure 8.

Note:

The configuration shown in Figure 8 assumes that the switches are not connected via Inter-Switch Link (ISL) in one fabric.

Viewing your current configuration

To view your current multipathing configuration, log on to your host machine as **root** and enter:

```
esxcfg-mpath -l
```

The output lists all known paths to each SAN LUN. An asterisk (*) indicates that the path is the current active path; # specifies the preferred path from the host machine to the LUN.

Tuning tips

- In an EVA array, HP recommends laying out your disks vertically. Distribute the disks across all enclosures, allowing the array to organize the RSS groups.
- In an environment with a mixture of SCSI and SATA drives, HP recommends storing your critical data on a dedicated RAID 1 SCSI disk group. You should configure SATA drives on their own disk group to be used for non-critical, less-accessed data.
- In a performance environment, HP recommends configuring one large disk group that can take advantage of all the disks in the array. If you need more than one disk group, HP recommends a maximum of three disk groups to optimize LUN manageability.
- HP recommends reviewing the OS-agnostic white paper, "HP StorageWorks 4000/6000/8000 Enterprise Virtual Array configuration best practices" (<http://h71028.www7.hp.com/ERC/downloads/4AA0-2787ENW.pdf>), which provides detailed information on EVA array functionality.
- For more information on optimizing EVA array performance, HP recommends reading the HP white paper, "Performance analysis of the HP StorageWorks Enterprise Virtual Array storage systems using HP StorageWorks Command View EVAPerf software tool," available at <http://h71028.www7.hp.com/ERC/downloads/5983-1674EN.pdf>.

Management and monitoring

- It is a best practice to use HP Command View Management 5.0 software to manage or monitor your array.
- HP Storage System Scripting Utility (SSSU) is supported by several guest operating systems (such as Linux, Microsoft Windows, and Novell NetWare) and can be used to verify array status from a command line interface on a remote Windows host.
- HP recommends involving a SAN administrator or storage specialist in the planning and implementation stages of the array deployment.
- Other tools (such as EVAperf, VirtualCenter, or the command-line utility **esxtop**) can be used to help IT staff understand the state of the host machine in a virtual infrastructure. By default, esxtop's display is updated every five seconds but update frequency can be changed using the **d** command-line option or the **s** interactive command.
- VMware Infrastructure 3 system resource variables (such as **Disk Statistics**) can be used to track physical disk usage, including total reads or writes per second and MB read and written per second. A display of physical disk information is provided per LUN per target per HBA. The display can be toggled using the interactive command, **d**. For more information, please refer to the VMware document, "Resource Management Guide" at http://www.vmware.com/pdf/vi3_esx_resource_mgmt.pdf.

XP128/1024 array

Connectivity

As with other arrays, a key best practice is to remain as current as possible with the latest tested and supported microcode. For HP StorageWorks XP128/1024 disk arrays, this implies the following:

- Minimum 21.14.18.00/00
- Recommended 21.14.18.00/00
- Mode 0x0C

Boot from SAN

- If a server is booting from SAN, a host group is required for the particular HBA.
- If multiple servers are booting from SAN and accessing the same physical port, a unique host group is required for each HBA.
- Each host group must be assigned to a port with a unique LUN 0.
- If LUN security is turned off for this port (for example, port CL1-A), all servers defined in the host groups are given access to all LUNs contained in the default host group.

Note:

For information on HP StorageWorks XP disk array microcode levels for specific guest operating systems and HBA and switched fabric configurations, contact your HP StorageWorks specialist or HP Services representative.

Tuning tips

- Balance your load across as many active components as reasonably possible.
- With multipathing, you can select either the Most-Recently Used (MRU) or fixed/preferred path. It is a best practice to use fixed paths, alternating the paths for each volume to minimize the risk of overrunning the array's CHIP-ports.
- Use dispersed Logical Unit Size Extension (LUSE)² in the back end of the LUN configuration. Try to balance the load across as many spindles as reasonably possible. Do not create a large number of very small volumes or a small number of very large volumes; HP recommends restricting volume size below 500 GB, if possible.
- Use Open-V LUN configuration, if possible.
- HP suggests implementing RAID 7+1 to stripe data across eight physical spindles – or even across two Array Control Processor (ACP) pairs – for any logical device. However, the high capacity per spindle available with RAID 7+1 may result in free capacity that cannot be used due to performance limits. If this is expected, RAID1 (2+2) or RAID-5 (3+1) may be a better choice.
- HP does not recommend the use of VMware Infrastructure 3-based spanned volumes. Spanning merely concatenates volumes and does not implement striping; consequently, from a performance perspective, there is little to be gained – apart from QDepth considerations. Spanning increases the complexity of setup and maintenance.

² A LUSE combines multiple logical devices (LDEVs). A dispersed LUSE includes LDEVs from multiple array groups.

- In general, for optimal performance there should be around two outstanding I/Os per physical spindle. With fewer I/Os, the array may be underutilized; with more I/Os, response times may begin to increase.
Other performance considerations include the type of load and the number of other (non-virtualized) servers sharing array groups with the host machines.
Assuming a mostly random I/O load, an appropriate formula might be:

$$\frac{[\text{Number of non-virtualized servers}] * [\text{Number of LUNs}] * [\text{QDepth}]}{2 * [\text{Number of physical spindles}]}$$
- Place swap files on internal disks for the following reasons:
 - Internal disks have better peak-response times
 - Information stored in swap files is not critical – you can afford to lose this information if the server crashes – but would create a significant extra load on the array.
- If you do implement a boot-from-SAN configuration, do not connect multiple storage arrays to either a single host machine – or a set of host machines. When booting, VMware Infrastructure 3 uses the first LUN on the first array detected.
The order the arrays report to the host machine is not deterministic and can change (for example, unplugging a cable from a switch port then plugging it back in may change the order of devices attached to the SAN).
Use zoning to hide other arrays connected to the SAN from host machines.
- HP recommends not selecting shared access mode. Unlike public access mode, shared access mode does not allow you to easily change stored metadata (to create new VMFS files for virtual disks, for example).
- RDMs are critical to giving you the ability to perform SAN-based local- and long-distance replication, including snapshots.
Note that RDM does not support the attachment of labels to volumes.
Additionally with RDM, the only way to identify a raw volume is through its hardware path, which can change (if, for example, the LUN mapping changes or if the SAN changes).

Management and monitoring

- The VMware **esxtop** utility can help IT staff understand the state of the host machine in a virtual infrastructure. By default, this utility's display is updated every five seconds; however, the update frequency can be changed using the **d** or **s** command.
- VMware Infrastructure 3 system resource variables (such as **Disk Statistics**) can be used to track physical disk usage, including total reads or writes per second and MB read and written per second. A display of physical disk information is provided per LUN per target per HBA. The display can be toggled using the **d** command.

More information

- HP suggests consulting with your HP StorageWorks specialist or HP Services representative to help you identify tools that can enhance your experience with HP StorageWorks arrays and VMware Infrastructure 3.
- More information on management and monitoring is provided in the VMware document, "Resource Management Guide" at http://www.vmware.com/pdf/vi3_esx_resource_mgmt.pdf.

Summary

This guide describes best practices for deploying HP StorageWorks storage arrays in a typical VMware Infrastructure 3 environment. It is not the intent of this guide to duplicate the strategies and best practices of other HP or VMware technical white papers.

The strategies and best practices offered here are presented at a very high level to provide general knowledge. Where appropriate, you are referred to specific white papers that provide more detailed information.

For more information

HP StorageWorks SAN Design Reference Guide	http://h18006.www1.hp.com/storage/saninfra/structure/index.html
HP StorageWorks Modular Smart Array (MSA)	http://www.hp.com/go/msa
HP StorageWorks Enterprise Virtual Array (EVA)	http://h18006.www1.hp.com/storage/arrayystems.html
HP StorageWorks XP arrays	http://h18006.www1.hp.com/storage/xparrays.html
HP StorageWorks Command View EVA	http://h18006.www1.hp.com/products/storage/software/cmdvieweva/index.html
HP StorageWorks Command View XP	http://www.hp.com/products1/storage/products/disk_arrays/xpstoragesw/commandview/index.html
Contact HP	Telephone numbers for worldwide technical support are listed on the HP support web site, http://welcome.hp.com/country/us/en/contact_us.html
VMware Hardware Compatibility List	http://www.vmware.com/pdf/vi3_san_guide.pdf
VMware SAN Configuration Guide	http://www.vmware.com/pdf/vi3_server_config.pdf
VMware Infrastructure 3 documentation	http://www.vmware.com/support/pubs/vi_pubs.html

© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. Linux is a U.S. registered trademark of Linus Torvalds.

August 2006 -4

