HITACHI
Inspire the Next

Hitachi Universal Storage Platform™ VM Best Practices for VMware Virtual Infrastructure

Best Practices Guide

By Patrick Flueckiger

March 2009

Summary

More and more organizations are adopting virtualization. Availability, balanced performance and ease of management are even more important than ever before with an increasing number of mission-critical applications being deployed in virtualized environments like VMware Virtual Infrastructure. The Hitachi Universal Storage Platform™ VM brings enterprise-class availability, performance and ease of management to organizations of all sizes. The best practice recommendations and configuration details in this document help to ensure successful planning, deployment and maintenance of Virtual Infrastructure environments that use Hitachi Universal Storage Platform VM storage systems.

Contributors

The information included in this document represents the expertise, feedback and suggestions of a number of skilled practitioners. The author would like to recognize and sincerely thank the following contributors and reviewers of this document:

- Alan Benway, Technical Operations
- Malcolm Brickwood, Product Marketing
- Elizabeth Graul, Application Solutions
- Dan Hood, Product Management
- James Hutchinson, Global Solutions Services
- Larry Korbus, Product Management
- Lisa Pampuch, Application Solutions
- Rob Simmons, Application Solutions

Table of Contents

Hitachi Universal Storage Platform™ VM Best Practices for VMware Virtual Infrastructure

Best Practices Guide

By Patrick Flueckiger

The many benefits of host virtualization — reducing cost of operation, increasing responsiveness, improving availability, enhancing performance, creating more environmentally friendly data centers and improving return on investment — mean that more and more organizations are adopting it. Availability, balanced performance and ease of management are even more important than ever before with an increasing number of mission-critical applications being deployed in virtualized environments like VMware Virtual Infrastructure.

The Hitachi Universal Storage Platform™ VM brings enterprise-class availability, performance and ease of management to organizations of all sizes that are dealing with an increasing number of virtualized business-critical applications. The Hitachi Universal Storage Platform VM with Hitachi Dynamic Provisioning software supports both internal and external virtualized storage, simplifies storage administration and improves performance to help reduce overall power and cooling costs.

The Hitachi Universal Storage Platform VM was the first external storage virtualization platform certified to work with VMware ESX Server. It provides end-to-end secure virtualization for VMware Virtual Infrastructure environments. With the ability to securely partition port, cache and disk resources, and to mask the complexity of a multivendor storage infrastructure, the Universal Storage Platform VM is an ideal complement to a VMware Virtual Infrastructure environment. With up to 1024 virtual ports for each physical Fibre Channel port, the Universal Storage Platform VM provides the connectivity to support large VMware server farms. Virtual Storage Machines can be created to align the host environment with the storage platform.

Hitachi Data Systems has the expertise to help you maximize availability, performance, and ease-of-management capabilities of the Universal Storage Platform VM in a Virtual Infrastructure environment. The best practice recommendations and configuration details in this document help to ensure successful planning, deployment and maintenance of Virtual Infrastructure environments that use Universal Storage Platform VM storage systems.

This white paper is intended for use by IT administrators who are in the following situations:

- Planning new deployments of VMware Virtual Infrastructure in an existing or new Universal Storage Platform VM storage environment
- Studying a move to server virtualization and looking for a storage platform to support their increased need for reliability and performance while reducing total cost of operation
- Adopting storage virtualization to support different data protection schemes
- Replacing older storage systems that require expensive maintenance or are not on VMware compatibility matrix
- Addressing concerns about storage availability and performance in a virtual server environment
- Seeking ways to reduce the cost and complexity of heterogeneous storage management

This paper is written for Virtual Infrastructure and IT administrators with solid working knowledge of Virtual Infrastructure and storage.

Hitachi Universal Storage Platform VM Overview

The Universal Storage Platform VM blends enterprise-class functionality with a smaller footprint to meet the business needs of entry-level enterprises and fast-growing mid-sized organizations, while supporting distributed or departmental applications in large enterprises.

Now, smaller organizations can enjoy the same benefits as large enterprises in deploying and managing their storage infrastructure in a way never possible before. The Universal Storage Platform VM, powered by the Hitachi Universal Star Network crossbar switch architecture, delivers proven and innovative controller-based virtualization, logical partitioning and replication for open systems and mainframe environments in a rack mounted storage services platform.

Figure 1. Hitachi Universal Storage Platform VM Virtual Storage Machines

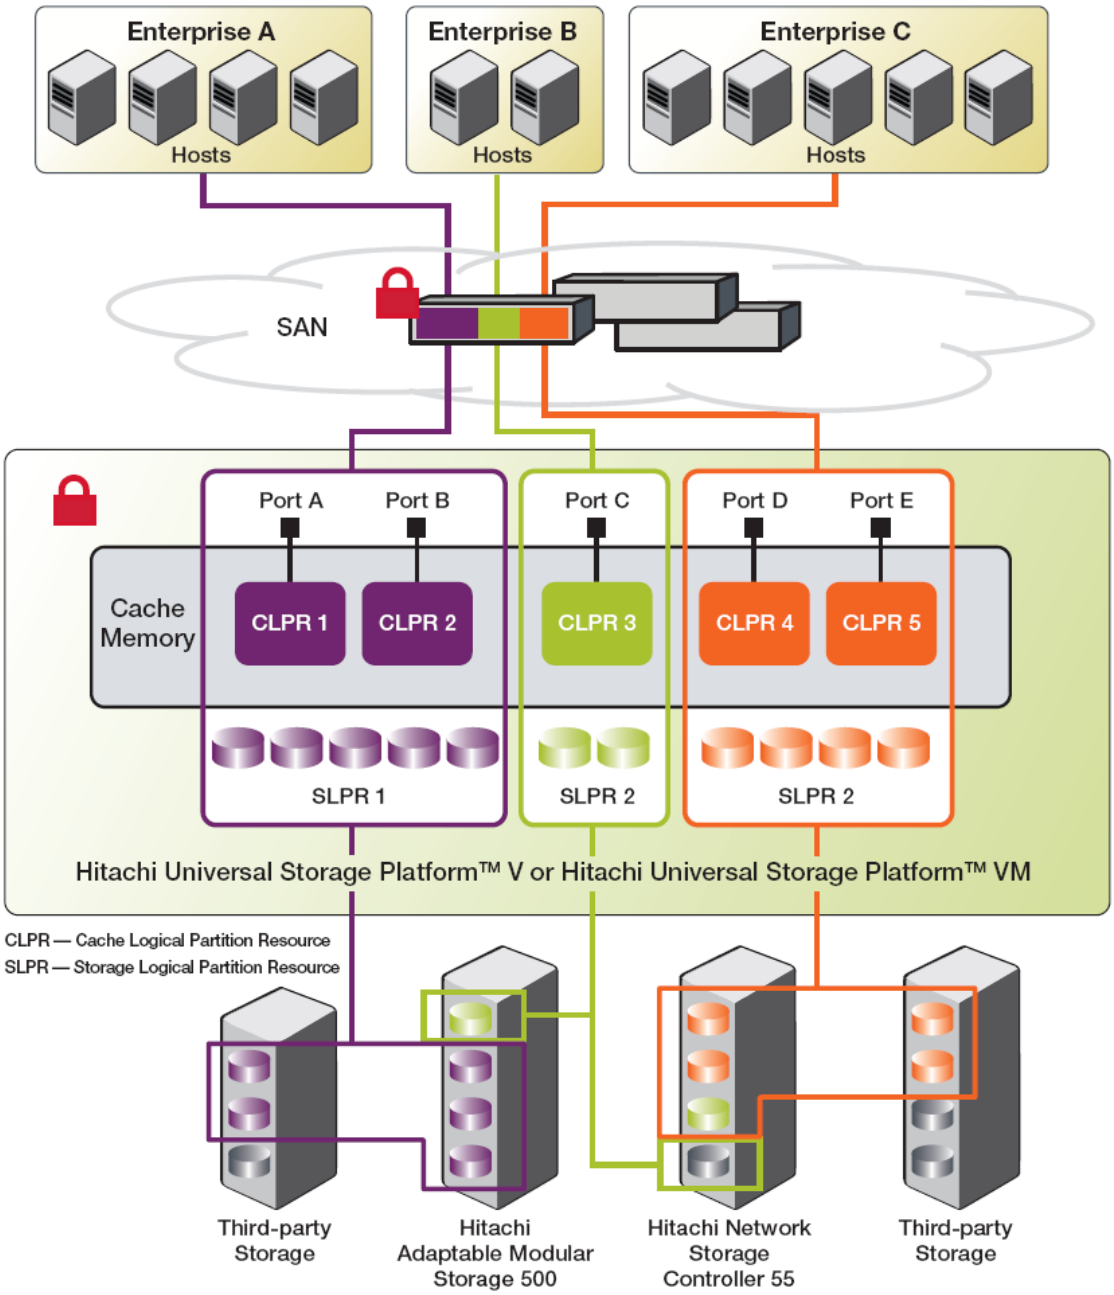

The industry's highest reliability and availability storage solution is backed by a set of storage and data services that include thin provisioning with Hitachi Dynamic Provisioning software, application centric storage management and logical partitioning, as well as simplified, unified data replication across heterogeneous storage systems.

The Hitachi Universal Storage Platform VM packages and delivers critical services like these:

- Virtualization of storage from Hitachi and other vendors into one pool
- Thin provisioning through Hitachi Dynamic Provisioning for nondisruptive volume expansion
- Security services, business continuity services and content management services
- Nondisruptive dynamic data migration from Hitachi and other storage systems

Features

These features of the Hitachi Universal Storage Platform VM let you deploy applications within a new framework, fully leverage and add value to current investments and more closely align IT to your business objectives:

- Excellent performance with 1.2 million I/Os 512 byte cached per second (IOPS), 13.3 GB/sec aggregate internal memory bandwidth
- Superior scalability providing up to 96PB of total storage capacity under management, including up to 105.7TB of Fibre Channel internal storage or 236.3TB of SATA internal storage
- Up to eight virtual storage machines with logical partitions, dedicating cache, ports and internally and externally attached capacity to ensure application quality of service (QoS)
- Superior multiplatform connectivity with up to 48 Fibre Channel, 24 IBM® ESCON® and 24 IBM FICON® host ports.
- Virtualization of externally attached storage from Hitachi, EMC, IBM, Sun, HP and other manufacturers
- Hitachi Dynamic Provisioning Software for thin provisioning of internal and virtualized external storage
- Nondisruptive movement, copy and migration of data between storage, including other vendors, without interrupting the application
- Replication across heterogeneous platforms for open and mainframe environments
- Multi-protocol consolidation with Fibre Channel, IBM FICON and ESCON and NAS
- Single-pane-of-glass management with Hitachi Storage Command Suite software
- Powered by the fourth generation Hitachi Universal Star Network crossbar switch architecture

Table 1 provides an overview of the features of the Hitachi Universal Storage Platform VM.

Table 1. Hitachi Universal Storage Platform VM Overview

Feature	*Hitachi Universal Storage Platform VM*
Aggregated IOPS	1.2 million cached 512 byte reads
Maximum internal capacity	240 disk drives
Maximum internal and external capacity	96PB
Maximum cache	128GB
Maximum control memory	16GB
Maximum Fibre Channel front-end interfaces	48 x 4 Gb/s Fibre Channel (FICON and ESCON also supported)
Maximum number of LUs	65,536
Maximum virtual host ports per physical port	1,024
Maximum virtual storage machines	8

Hitachi Virtual Partition Manager Software

Hitachi Virtual Partition Manager software logically partitions Universal Storage Platform VM cache, ports and disk capacity, including capacity on externally attached storage systems. The software enables administrators to create Hitachi Virtual Storage Machines (formerly called Private Virtual Storage Machines). Each machine is an isolated group of storage resources, with their own storage partition administrator. Logical partitioning guarantees data privacy and QoS for host virtualized and non-virtualized environments sharing the same storage platform.

Hitachi Universal Volume Manager Software

Universal Volume Manager software simplifies migration of datastores to the Universal Storage Platform VM. The software eliminates the need to reformat LUs or perform complex extent mapping, providing a minimally disruptive migration process. Once virtualized behind the Universal Storage Platform VM, external LUs gain access to all Hitachi storage services.

Hitachi Dynamic Provisioning Software

Hitachi Dynamic Provisioning software provides the Universal Storage Platform VM with thin provisioning services. Thin provisioning gives applications access to virtual storage capacity. Applications accessing virtual, thin provisioned volumes are automatically allocated physical disk space, by the storage system, as they write data. This means volumes use enough physical capacity to hold application data, and no more. All thin provisioned volumes share a common pool of physical disk capacity. Unused capacity in the pool is available to any application using thin provisioned volumes. This eliminates the waste of overallocated and underutilized storage.

Dynamic Provisioning software also simplifies storage provisioning and automates data placement on disk for optimal performance. Administrators do not need to micro-manage application storage allocations or perform complex, manual performance tuning. In addition, physical storage resources can be added to the thin provisioning pool at any time, without application downtime.

Hitachi Global Solution Services

Complex Virtual Infrastructure environments need a full analysis to determine the specific performance and capacity requirements of the storage systems. Hitachi Data Systems Global Solution Services helps you maximize investments and efficiencies with a suite of professional services. For the Hitachi Universal Storage Platform VM, Global Solution Services can assist with the following important tasks:

- Remote copy planning and design

- Implementation of Hitachi TrueCopy® Synchronous and Hitachi Universal Replicator
- Implementation of Hitachi ShadowImage® Replication software
- Data migration planning, design and implementation

For more information about services to help meet regulatory compliance requirements, protect data, reduce total cost of ownership or develop a disaster recovery plan, contact your Hitachi Data Systems representative or visit www.hds.com.

Basic Storage System Setup

Follow the best practices provided in this section to successfully deploy VMware Virtual Infrastructure on Hitachi Universal Storage Platform VM systems. The information provided might apply directly to your environment or might need some modifications. Always take into account specific requirements from your environment and check with a Hitachi Data Systems representative.

Host Groups per HBA Versus Host Groups per ESX Hosts or VMware Cluster

To present a set of common, shared LUs to multiple ESX hosts or to a VMware cluster, host groups can be created either per HBA port (that is, per WWPN) or per a group of ESX hosts or VMware cluster.

A host group created on an HBA port basis contains the HBA's WWPN and a set of common, shared LUs (that is, only one WWPN, multiple LUs). A host group created per group of ESX hosts or per VMware cluster contains at least one WWPN from every ESX host and multiple LUs (that is, multiple WWPNs, multiple LUs). Every LU must be presented with the same host LU ID to every host or VMware treats the LU as a snapshot LU and disables access to the VMFS by default.

Although both concepts are supported, Hitachi Data Systems recommends creating host groups per HBA port (that is, per WWPN).

Figure 2. Host Groups Created per HBA Port

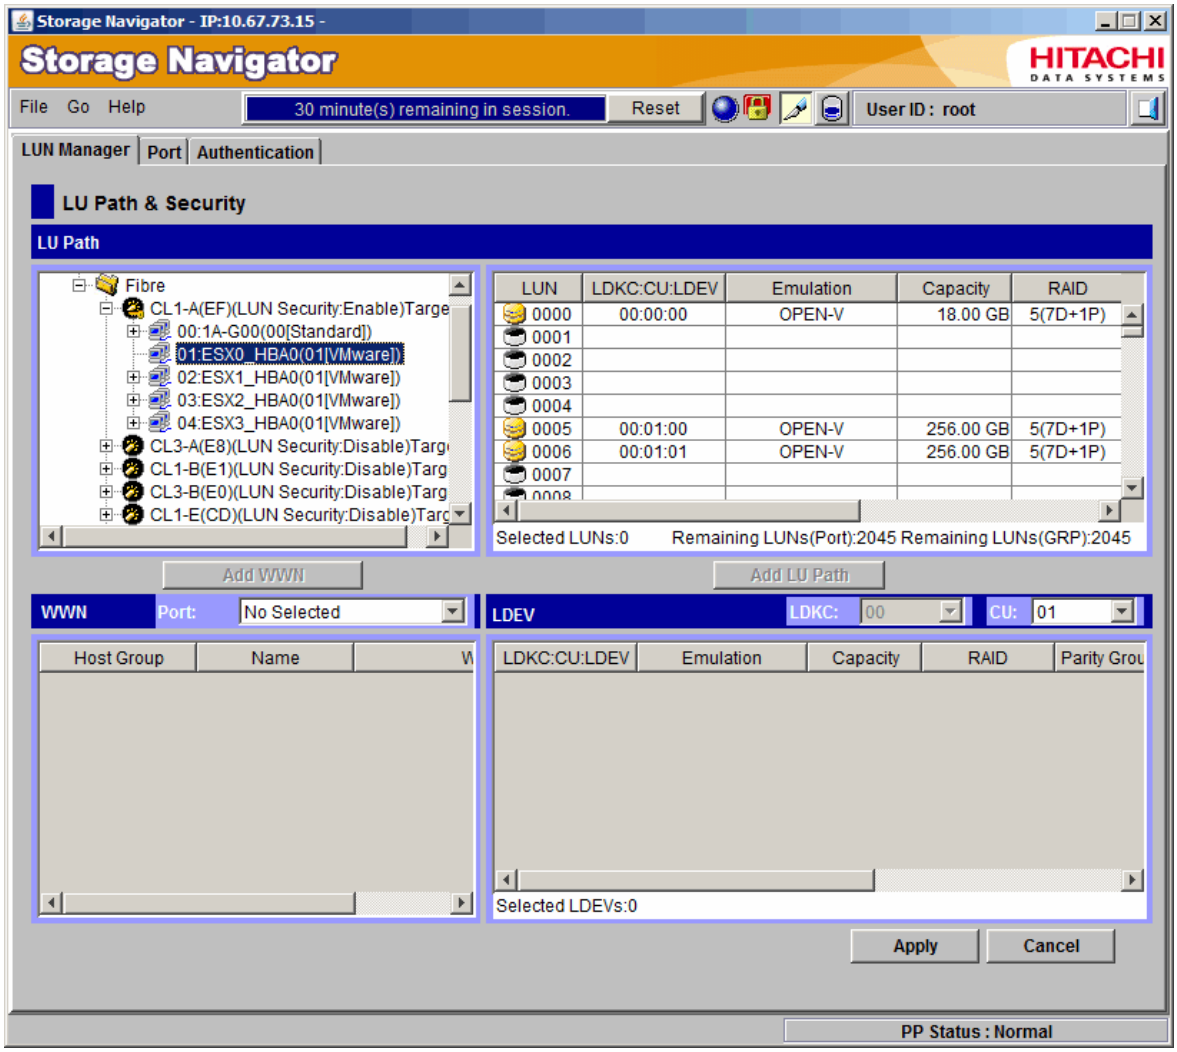

Table 2 lists some advantages and disadvantages of creating host groups per HBA port versus host groups per ESX hosts or VMware cluster.

Table 2. Hitachi Data Systems Host Group Recommendations

Host Group Configuration	*Advantages*	*Disadvantages*
One host group per HBA (HBA port)	Allows boot from SAN. Better flexibility. Changes to the host group affect only one host. Better troubleshooting possibilities and log analysis.	More host groups to be created and maintained.
One host group per ESX hosts or VMware cluster	Fewer host groups to be created and maintained.	Does not allow ESX hosts to boot from SAN. Less flexibility (for example, an additional LU cannot be presented to only one ESX host). Changes to the host group might affect all hosts.

Host Modes

To create host groups for ESX hosts, choose **01[VMware]** or **21[VMware extension]** from the **Host Mode** drop-down menu. Host Mode 21[VMware extension] allows the storage administrator to expand a LU using Logical Unit Size Expansion (LUSE) while the LU is mapped to the host.

Figure 3. Host Mode

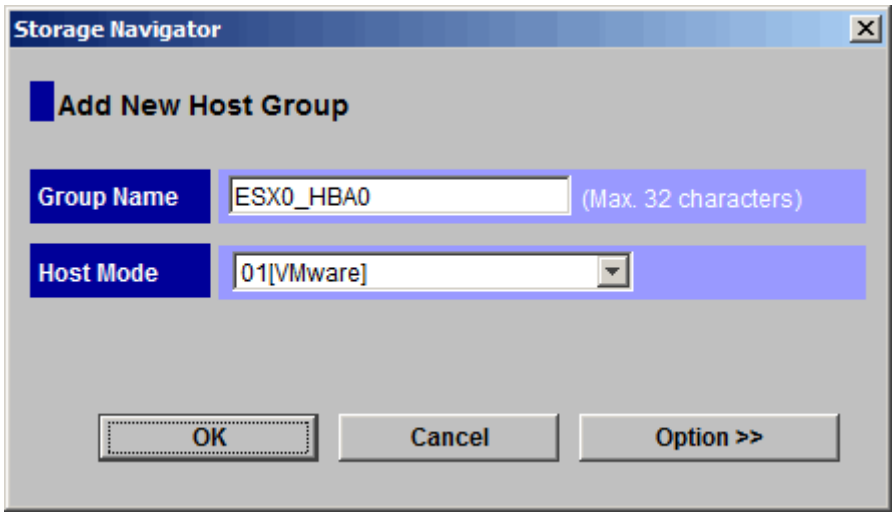

Storage Provisioning

Capacity and performance cannot be considered independently. Performance always depends on and affects capacity and vice versa. That's why it's very difficult or impossible in real-life scenarios to provide best practices for the best LU size, the number of virtual machines that can run on a single VMFS and so on without knowing both capacity and performance requirements. However, several factors must be considered when planning storage provisioning for a Virtual Infrastructure environment.

Size of LU

When determining the right LU size, consider the factors listed in Table 3. These factors are especially important from a storage system perspective. In addition, the individual virtual machine's capacity and

performance requirements (basic virtual disk requirements, virtual machine swap space, spare capacity for virtual machine snapshots and so on) must also be considered.

Table 3. LU Size Considerations

Factor	*Comment*
Minimum LU size	The smallest LU that can be created on the Hitachi Universal Storage Platform VM is 96,000 blocks (approx. 46 MB). The smallest VMFS that can be created on an ESX host is approx. 1GB.
Maximum LU size	The largest internal LU that can be created on the Hitachi Universal Storage Platform VM is 2.999TB. The largest VMFS that can be created on an ESX host is 2TB (without using extents, 64TB when using extents).
Disk queues	More LUs offer more disk queues and therefore better performance.
SCSI reservations	Because SCSI reservations are set per LU, more but smaller LUs typically lead to less SCSI reservation conflicts and therefore better performance.
Volume migration	Smaller LUs can be migrated using Hitachi Tiered Storage Manager to a broader range of possible target RAID groups.
Data replication	Using more but smaller LUs offers better flexibility and granularity when replication a VMFS within a storage system (ShadowImage Replication software, Hitachi Copy-on-Write Snapshot software) or across storage systems (TrueCopy Synchronous or Hitachi Universal Replicator).
VMware VMFS upgrades	Using more but smaller LUs offers better flexibility and granularity when upgrading a VMFS.
VMware vCenter Site Recovery Manager	Using more but smaller LUs offers better flexibility and granularity when creating protection groups and recovery plans in VMware vCenter Site Recovery Manger.

Considering all of the factors in Table 3, a good LU size to start with is 256GB or 512GB. This number might vary from environment to environment.

Number of Virtual Machines per VMFS, per RAID Group

The number of virtual machines that can run simultaneously on a VMFS depends on the aggregated capacity and performance requirements of the virtual machines. Because all LUs on a particular RAID group share the performance and capacity offered by the RAID group, Hitachi Data Systems recommends dedicating RAID groups to an ESX host or a group of ESX hosts (for example, an ESX cluster) and not assigning LUs from the same RAID group to the Virtual Infrastructure environment and other non-ESX hosts. This prevents the ESX I/O from affecting or being affected by other applications and LUs on the same RAID group and makes management cleaner.

Follow these best practices:

- Create and dedicate RAID groups to your ESX hosts (that is, VMware cluster)
- Create LUs of 256GB or 512GB in size, or adjust the LU size considering the factors listed in Table 3. Always present LUs with the same LUN if they are shared with multiple hosts.
- Create VMFSs on the LUs as needed. Consider choosing a datastore name with a reference to the internal LU ID.
- Create and run virtual machines on the VMFS files.
- Monitor and measure the capacity and performance usage of the RAID group with VMware vCenter Server, Hitachi Tuning Manager software and Hitachi Performance Monitor software.

Monitoring and measuring the capacity and performance usage of the RAID group results in one the following cases:

- If all of the capacity offered by the RAID group is used but performance of the RAID group is still good, add RAID groups and therefore more capacity. In this case, consider migrating the LUs to a different RAID group with less performance using Hitachi Tiered Storage Manager.
- If all of the performance offered by the RAID group is used but capacity is still available, do not use the remaining capacity by creating more LUs and VMFSs because this leads to even more competition on the RAID group and overall performance for the virtual machines residing on this RAID group is affected. In this case, leave the capacity unused and add more RAID groups and therefore more performance resources. Also consider migrating the LUs to a different RAID group with better performance.

In a real environment, it is not possible to use 100 percent of both capacity and performance of a RAID group, but the usage ratio can be optimized by actively monitoring the systems and moving data to the appropriate storage tier if needed using Hitachi Tiered Storage Manager. An automated solution using these applications from the Hitachi Storage Command Suite helps to reduce the administrative overhead and optimize storage utilization.

Fibre Channel Front-end Ports

Provisioning storage on two Fibre Channel front-end ports (on one port per storage cluster) is sufficient for redundancy on the Hitachi Universal Storage Platform VM. This results in two paths to each LU from an ESX host's point of view. More Fibre Channel front-end ports can be selected for availability and business continuity reasons.

Basic ESX Host Setup

Storage in a VMware Virtual Infrastructure environment is typically provisioned to a group of ESX hosts — to an ESX cluster — sharing a common set of LUs. Too many SCSI reservation conflicts or improper settings for queue depths can lead to performance issues.

Path Policies

Hitachi recommends using a fixed path policy for all of its storage systems. Table 4 describes the advantages and disadvantages of various multipathing policies.

Table 4. Virtual Infrastructure Multipathing Policies

Multipathing Policy	*HDS Recommendation*	*Advantage*	*Disadvantage*
Most recently used	Not recommended	None.	Manually set path policy for every LU on every ESX host. No control over w rkload distribution across SAN fabrics and storage ports.
Fixed	Recommended	This is the default path policy for VMware ESX 3.5 for any Hitachi storage system. No manual changes needed.	Workload must still be balanced across host bus adapters and SAN fabrics.
Round robin	Experimental support only from VMware	Workload is automatically balanced across host bus adapters and SAN fabrics.	Manually set path policy for every LU on every ESX host.

Enterprise symmetric active-active controller architecture and Virtual Infrastructure's round robin multipathing policy dramatically simplifies the setup and management of a Virtual Infrastructure environment because the workload is automatically balanced across the following:

- Host bus adapters and SAN fabrics (due to multipathing policy round robin)
- Storage systems front-end Fibre Channel ports (due to multipathing policy round robin and symmetric active-active controller architecture)

SCSI Reservations

Many operations — such as virtual machine power on or off, VMotion, virtual machines running with virtual disk snapshots or growing a file — require getting a file lock or a metadata lock in VMFS, resulting in a temporary SCSI reservation. During this short period of time, the entire LU is exclusively reserved and locked by a host and is therefore not accessible by other hosts. Excessive SCSI reservations by one or multiple ESX hosts cause SCSI reservation conflicts and can cause performance degradation on other ESX hosts accessing the same VMFS or LU.

You can search VMkernel log files for SCSI reservation conflict warnings. For more information about how to reduce SCSI reservations and how to avoid SCSI reservation conflicts, see VMware's Fibre Channel SAN Configuration Guide.

From a storage provisioning point of view, SCSI reservation conflicts can be reduced by creating more but smaller LUs to achieve a given capacity.

Logical Unit Size Expansion

Logical Unit Size Expansion (LUSE) is a concatenation of LUs and not a performance feature. Hitachi Data Systems does not recommend extensively using LUSE to provision storage to a VMware Virtual Infrastructure

environment due to the increased time used to reserve and release SCSI reservations compared to a non-LUSE LU. With an increased number of LUSE members and ESX hosts, sharing a LUSE LU can increase the amount of SCSI reservation conflicts and can lead to performance issues.

Queue Depth

The Hitachi Universal Storage Platform VM supports a maximum queue depth per Fibre Channel front-end port of 4,096 using eight port FED features or 2,048 per Fibre Channel front-end port when using 16 port FED features. Each LU has a nominal queue depth limit of 32. This value determines the maximum queue depth setting on the VMware ESX hosts. A large queue depth does not provide better performance per se. This can cause the response times to rapidly increase while gaining a small increase in achieved IOPS at the host. Instead, monitor the queued commands on the hosts and adjust the HBA queue depth to ensure that the front-end port does not exceed the maximum queue depth of 2,048 or 4,096 per Fibre Channel port. Use the `esxtop` command to determine the number of commands queued per ESX host.

Table 5 lists supported maximum queue depth settings.

Table 5. Queue Depth Settings for Hitachi Universal Storage Platform VM

Configuration	*Supported Maximum Setting*
Path Policy Fixed or Round Robin, Host Groups created per HBA (per host WWPN)	Maximum queue depth of 32 per LU Maximum aggregated queue depth of 4,096 per Fibre Channel port using eight port FED features. Maximum aggregated queue depth of 2,048 per Fibre Channel port using 16 port FED features.
Path Policy Fixed or Round Robin, Host Groups created per ESX cluster	Maximum queue depth 32 divided by the number of ESX hosts (active HBAs) with access to the LU Example: 32/4=8 (4 ESX hosts, queue depth of 8 on each ESX hosts HBA) Maximum aggregated queue depth of 4,096 per Fibre Channel port using eight port FED features. Maximum aggregated queue depth of 2,048 per Fibre Channel port using 16 port FED features.

VMware Snapshot LUs

At the moment that an LU is formatted with the VMFS by an ESX host, it writes a signature that contains the following information to the disk:

- Serial number of the storage system
- Internal LU ID (LDKC:CU:LDEV)
- External LU ID (host LUN)

Whenever an ESX host has access to a VMFS, it reads and compares the signature with its current view. If it does not match, the host disables access to the VMFS (not the LU) by default and generates the following warning:

```
ALERT: LVM: 4941: vmhba0:0:1:1 may be a snapshot: disabling access. See
resignaturing section in the SAN config guide.
```

In this case, access to a LU can be enabled with one of the following ESX advanced parameters:

- `LVM.EnableResignature`
- `LVM.DisableSnapshotLUN`

The default values of these parameters disable access to the VMFS.

Table 6 lists the situations where an ESX host might treat a VMFS as a snapshot LUN. In these situations, the host lists and recognizes the LU in the Storage Adapter screen of Virtual Infrastructure Client. A SAN rescan on the HBA might be necessary to discover the LU if it is newly mapped. The VMFS might appear in the Storage screen of Virtual Infrastructure client depending on the situation and the settings of the advanced ESX parameters.

Table 6. ESX Host Treats VMFS as Snapshot LU Situations

Situation	*ESX Behavior*
In-system data replication (ShadowImage, Copy-on-Write)	The LU is treated as a snapshot and access to the VMFS is disabled by default. Reason: At least the internal LU ID changed. It is possible that the external LU ID also changed.
Cross-system data replication (TrueCopy Synchronous, Hitachi Universal Replicator)	The LU is treated as a snapshot and access to the VMFS is disabled by default. Reason: At least the serial number of the storage system changed. Note: VMware vCenter Site Recovery Manager uses cross-system data replication and handles the situation.
Present LU to multiple ESX hosts with different host LUNs	The LU is treated as a snapshot by some or all ESX hosts and access to the VMFS is disabled by default. Reason: The external LU ID (host LUN) has changed for some or all ESX hosts.
Volume migration	The LU is not treated as a snapshot, the VMFS is accessible. Reason: Volume migration is transparent and non-disruptive to the host. Neither serial number of the storage system, internal LU ID nor external LU ID (host LUN) changed.
Storage virtualization	For the process only of virtualizing storage (migrating from non-virtualized storage to virtualized storage) the LU is treated as a snapshot and access to the VMFS is disabled by default. Reason: At least the serial number of the storage system and the internal LU ID changed.
Online LUSE	Does not affect the VMFS.

Storage Virtualization

The Hitachi Universal Storage Platform VM offers built-in storage virtualization that allows other storage systems (from Hitachi and from third parties) to be attached (or virtualized) behind the Hitachi Universal Storage Platform VM. From an ESX host's point of view, virtualized storage is accessed through the Hitachi Universal Storage Platform VM and appears like internal, native storage capacity. The virtualized storage systems immediately inherit every feature available on the Hitachi Universal Storage Platform VM (data replication, Hitachi Dynamic Provisioning, and so on) and thus enable management and replication using Hitachi software.

To virtualize storage systems behind a Hitachi Universal Storage Platform VM for a VMware Virtual Infrastructure environment, use the following high-level checklist. Although the process itself is conceptually simple and usually only requires logical reconfiguration tasks, always check and plan this process with your Hitachi Data Systems representative.

1. Quiesce I/O to and unmap the LUs from the ESX hosts on the storage system to be virtualized.
2. Reconfigure the SAN zoning as needed.
3. Map the LUs to the Hitachi Universal Storage Platform VM using the management tools available on the storage system to be virtualized.
4. Map the (virtualized) LUs to the ESX hosts.

   The ESX hosts treat the LUs as snapshot LUs because the signature on the volumes does not match their current view. For more information, see the "VMware Snapshot LUs" section in this document. Before renaming snapshot LUs back to their original names, every reference to the old datastore name within VMware vCenter must be removed. This can be achieved by removing every virtual machine from inventory that has its virtual disks on the particular datastore and then removing the old datastore itself.

Hitachi Dynamic Provisioning

Storage can be provisioned to a VMware Virtual Infrastructure environment using Hitachi Dynamic Provisioning. Virtual DP volumes have a defined size, viewed by the ESX hosts as any other normal volume and initially do not allocate any physical storage capacity from the HDP pool volumes. Data is written and striped across the HDP pool volumes in a fixed size that is optimized to achieve both performance and storage area savings.

Figure 4. Hitachi Dynamic Provisioning Concept Overview

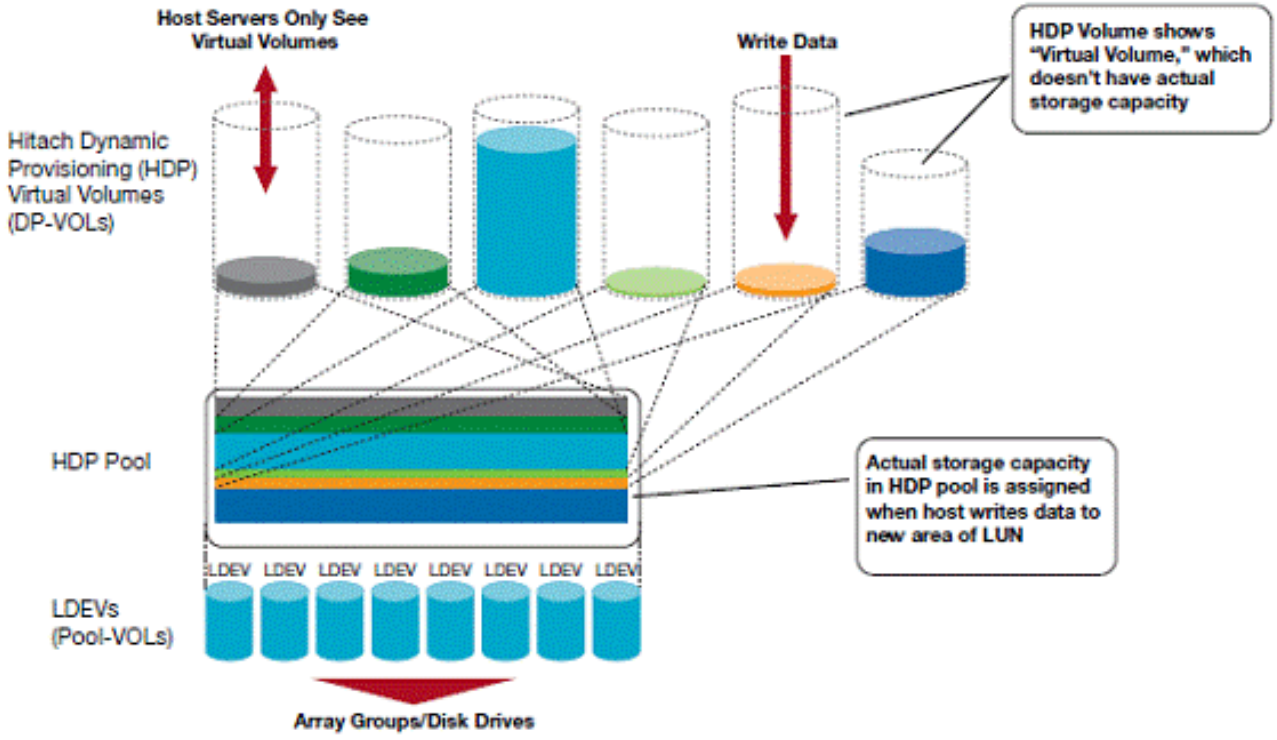

Table 7 describes the effects of configuration steps on HDP pool volume capacity allocation.

Table 7. Hitachi Dynamic Provisioning Capacity Allocation

Configuration Step	*Effect on HDP Pool Volume Capacity Allocation*
Map an HDP volume to ESX hosts.	This process does not allocate any physical capacity on the HDP pool volumes.
Create a VMFS on the HDP volume.	This process does allocate some physical capacity on the HDP pool volumes to write VMFS metadata.
Create a virtual machine on the VMFS and create the virtual machine's disk file (`.vmdk` file).	This process does not allocate additional capacity on the HDP pool volumes. The `.vmdk` file is created but its capacity is not allocated.
Install an operating system in the virtual machine.	This process does allocate capacity on the HDP pool volumes depending on the file system being used in the virtual machine and the amount of data written to the virtual machine's disk file (`.vmdk` file).
Deploy a virtual machine from a template.	This process does allocate the whole capacity of the virtual machine's disk file on the HDP pool volumes.
Delete data within the virtual machine.	The capacity remains allocated on the HDP pool volumes but might be reused by the virtual machine.
Delete the virtual machine and delete the virtual machine's disk file (`.vmdk` file)	The capacity remains allocated on the HDP pool volumes but might be reused by the ESX host.

Hitachi Data Systems Corporation

Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.uk@hds.com

AS-007-00 March 2009