

Optimize Hitachi Storage and Server Platforms in VMware vSphere 5.5 Environments

Best Practices Guide

By Daniel Worden

February 2016

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Contents

Solution Components	3
Hitachi Virtual Storage Platform G1000	3
Hitachi Virtual Storage Platform	3
Hitachi Unified Storage VM	4
Hitachi Unified Storage	4
Hitachi Accelerated Flash	4
Hitachi NAS Platform	5
Hitachi Dynamic Provisioning	5
Hitachi Dynamic Tiering	6
Hitachi Unified Compute Platform Director	6
Hitachi Compute Blade 500	7
Hitachi Compute Blade 2000	7
Hitachi Compute Rack 210H	7
Hitachi Compute Rack 220H	8
Hitachi Data Instance Manager	8
Hitachi Data Protection Suite, Powered by CommVault, with Hitachi Thin Image	9
Hitachi Storage Plug-Ins for VMware vCenter	10
Hitachi Compute Plug-in for VMware vCenter	10
VMware vSphere 5.5	10
Hitachi Storage Provider for VMware vCenter, and VMware vSphere API for Storage Awareness ..	11
VMware vStorage APIs for Array Integration	11
Performance and Scalability Best Practices	12
VMware vSphere Storage APIs for Array Integration	12
VMware vSphere Storage I/O Control	13
VMware vSphere Storage DRS	16
Disk Alignment	16
VMDK Provisioning Type	17
Hitachi Accelerated Flash	17
Reliability and Scalability Best Practices	19
Multipathing	19
Redundant HBA	20
Redundant Fibre Channel Switches	20
Zoning	20
Host Group Configuration	21

Contents

Usability Best Practices	22
Storage Manager for vCenter.....	22
Hitachi Compute Plug-in for VMware vCenter.....	23
VMware vStorage API for Storage Awareness.....	25
Hitachi Dynamic Tiering	27
Virtual Desktop Infrastructure Environment Best Practices.....	28
VMware Horizon View	31
Citrix XenDesktop Linked Clones	31
Flex SE VMDKs and Flex SE Linked Clones.....	31
VDI User Data Redirection	32
Resilience and Recoverability Best Practices.....	33
vSphere Metro Storage Cluster	33
VMware Site Recovery Manager.....	33
Hitachi Data Protection Suite, Powered by ComVault with Hitachi Thin Image	34
Hitachi Data Instance Manager.....	34
Conclusion.....	35

Optimize Hitachi Storage and Server Platforms in VMware vSphere 5.5 Environments

Best Practices Guide

Today's demanding economic and IT environment makes it more important than ever that IT administrators provide high-performance, scalable, highly available and easy-to-manage infrastructures to maximize both business and IT efficiency. With this goal in mind, organizations are expanding their use of virtualization to more business-critical and mission-critical (tier one and tier two) applications. As they do, the storage, servers, and software deployed to support this environment becomes increasingly important.

Hitachi Data Systems industry-differentiated storage and server architecture make Hitachi solutions the ideal solution for VMware vSphere environments. The depth and variety of Hitachi Data Systems storage and server products ensures that customers are able to build scalable, reliable, and resilient server and VDI solutions that are easily usable.

Some of the Hitachi Data Systems storage and server products covered in this best practices paper are listed in Table 1.

Table 1. Hitachi Data Systems Storage and Server Products

Hardware	Product
Storage	Hitachi Virtual Storage Platform G1000 (VSP G1000) Hitachi Virtual Storage Platform (VSP) Hitachi Unified Storage VM (HUS VM) Hitachi Unified Storage 150 (HUS) Hitachi Accelerated Flash (HAF)
Servers	Hitachi Compute Blade 500 (CB 500) Hitachi Compute Blade 2000 (CB 2000) Hitachi Compute Rack 210H (CR 210H)

Some of the Hitachi Data Systems software plugin products covered in this best practices paper are listed in Table 2.

Table 2. Hitachi Data Systems Plugin and Adapter Products

Software	Product
Hitachi Storage Software	Hitachi Dynamic Provisioning Hitachi Dynamic Tiering vStorage API for Array Integration (VAAI)
Hitachi Adapters and Plugins for VMware vSphere	Storage Provider for VMware vStorage API for Storage Awareness (VASA) Hitachi Dynamic Link Manager (HDLM) for VMware vSphere Hitachi Storage Replication Adapter for VMware vCenter Site Recovery Manager (SRA) Storage Manager for VMware vCenter Server Manager for VMware vCenter
Hitachi Integrations with VMware vSphere	Hitachi Data Instance Director Hitachi Data Protection Suite, powered by CommVault, with Hitachi Thin Image

Switching is always a key component in any virtualized infrastructure. Hitachi Data Systems designs solutions using both Cisco and Brocade Ethernet devices and Brocade Fiber Channel devices. The best practices covered in the document apply to both vendors.

Solution Components

This section describes the hardware and software components mentioned in this best practices guide.

Hitachi Virtual Storage Platform G1000

Hitachi Virtual Storage Platform G1000 provides an always-available, agile, and automated foundation that you need for a continuous infrastructure cloud. This delivers enterprise-ready software-defined storage, advanced global storage virtualization, and powerful storage.

Supporting always-on operations, Virtual Storage Platform G1000 includes self-service, non-disruptive migration and active-active storage clustering for zero recovery time objectives. Automate your operations with self-optimizing, policy-driven management.

Virtual Storage Platform G1000 supports Oracle RAC and VMware Metro Storage Cluster.

Hitachi Virtual Storage Platform

[Hitachi Virtual Storage Platform](#) is the first 3-D scaling storage platform designed for all data types. Its storage architecture flexibly adapts for performance, capacity, and multi-vendor storage. Combined with the unique Hitachi Command Suite management platform, it transforms the data center.

- **Scale Up** — Meet increasing demands by dynamically adding processors, connectivity, and capacity in a single unit. Provide the highest performance for both open and mainframe environments.
- **Scale Out** — Meet multiple demands by dynamically combining multiple units into a single logical system with shared resources. Support increased demand in virtualized server environments. Ensure safe multi-tenancy and quality of service through partitioning of cache and ports.
- **Scale Deep** — Extend storage value by virtualizing new and existing external storage systems dynamically. Extend the advanced functions of Hitachi Virtual Storage Platform to multivendor storage. Offload less demanding data to external tiers to save costs and to optimize the availability of tier-one resources.

The first 3-D scaling storage platform designed for all data types that flexibly adapts for performance, capacity, and multi-vendor storage with increased scalability using flash media or disk drives in a large pool

Hitachi Unified Storage VM

[Hitachi Unified Storage VM](#) is an entry-level enterprise storage platform. It combines storage virtualization services with unified block, file, and object data management. This versatile, scalable platform offers a storage virtualization system to provide central storage services to existing storage assets.

Unified management delivers end-to-end central storage management of all virtualized internal and external storage on Unified Storage VM. A unique, hardware-accelerated, object-based file system supports intelligent file tiering and migration, as well as virtual NAS functionality, without compromising performance or scalability.

The benefits of Unified Storage VM are the following:

- Enables the move to a new storage platform with less effort and cost when compared to the industry average
- Increases performance and lowers operating cost with automated data placement
- Supports scalable management for growing and complex storage environment while using fewer resources
- Achieves better power efficiency and with more storage capacity for more sustainable data centers
- Lowers operational risk and data loss exposure with data resilience solutions
- Consolidates management with end-to-end virtualization to prevent virtual server sprawl

Hitachi Unified Storage

Hitachi Unified Storage is a midrange storage platform for all data. It helps businesses meet their service level agreements for availability, performance, and data protection.

The performance provided by Hitachi Unified Storage is reliable, scalable, and available for block and file data. Unified Storage is simple to manage, optimized for critical business applications, and efficient.

Using Unified Storage requires a smaller capital investment. Deploy this storage, which grows to meet expanding requirements and service level agreements, for critical business applications. Simplify your operations with integrated set-up and management for a quicker time to value.

Unified Storage enables extensive cost savings through file and block consolidation. Build a cloud infrastructure at your own pace to deliver your services.

Hitachi Unified Storage 150 provides reliable, flexible, scalable, and cost-effective modular storage. Its symmetric active-active controllers provide input-output load balancing that is integrated, automated, and hardware-based.

Both controllers in Unified Storage 150 dynamically and automatically assign the access paths from the controller to a logical unit (LU). All LUs are accessible, regardless of the physical port or the server that requests access.

Hitachi Accelerated Flash

[Hitachi Accelerated Flash](#) features a flash module built specifically for enterprise-class workloads. Developed for [Hitachi Virtual Storage Platform](#), Accelerated Flash is available for [Hitachi Unified Storage VM](#).

Accelerated Flash features innovative Hitachi-developed embedded flash memory controller technology. Hitachi flash acceleration software speeds I/O processing to increase flash device throughput.

Hitachi NAS Platform

[Hitachi NAS Platform](#) is an advanced and integrated network attached storage (NAS) solution. It provides a powerful tool for file sharing, file server consolidation, data protection, and business-critical NAS workloads.

- Powerful hardware-accelerated file system with multiprotocol file services, dynamic provisioning, intelligent tiering, virtualization, and cloud infrastructure
- Seamless integration with Hitachi SAN storage, [Hitachi Command Suite](#), and [Hitachi Data Discovery Suite](#) for advanced search and index
- Integration with [Hitachi Content Platform](#) for active archiving, regulatory compliance, and large object storage for cloud infrastructure
- **Hitachi NAS Platform** — Network-attached storage solution used for file sharing, file server consolidation, data protection, and business-critical NAS workloads

For recommendations on best practices when using Hitachi NAS Platform in a VMware vSphere environment please see [Hitachi NAS Platform Best Practices Guide for NFS with VMware vSphere](#).

Hitachi Dynamic Provisioning

On Hitachi storage systems, [Hitachi Dynamic Provisioning](#) provides wide striping and thin provisioning functionalities.

Using Dynamic Provisioning is like using a host-based logical volume manager (LVM), but without incurring host processing overhead. It provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) of a logical size you specify of up to 60 TB created against it without allocating any physical space initially.

Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. Dynamic Provisioning distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Unified Storage Platform VM

When used with [Hitachi Unified Storage VM](#), Hitachi Dynamic Provisioning has the benefit of thin provisioning. Physical space assignment from the pool to the dynamic provisioning volume happens as needed using 42 MB pages, up to the logical size specified for each dynamic provisioning volume. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. You can rebalance an expanded pool across the current and newly added RAID groups for an even striping of the data and the workload.

Hitachi Unified Storage

When used with [Hitachi Unified Storage](#), Hitachi Dynamic Provisioning has the benefit of thin provisioning. Physical space assignment from the pool to the dynamic provisioning volume happens as needed using 1 GB chunks, up to the logical size specified for each dynamic provisioning volume. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. You can rebalance an expanded pool across the current and newly added RAID groups for an even striping of the data and the workload.

Hitachi Dynamic Tiering

[Hitachi Dynamic Tiering](#) eliminates manual data classification and movement between storage tiers. This optimizes tiered storage usage while improving performance.

Instead of manually provisioning space from several storage technologies with different performance and cost characteristics, Hitachi Dynamic Tiering enables the management of multiple storage tiers as a single entity. By leveraging the thin provisioning and wide striping features of Hitachi Dynamic Provisioning, Hitachi Dynamic Tiering presents a virtual volume with embedded smart tiering. It monitors access and moves data at the 42 MB page level.

Breaking the volume into pages, Hitachi Dynamic Tiering automatically moves infrequently referenced pages to lower cost tiers of storage. Moving pages instead of entire data sets or files reduces the time and storage space required to migrate data.

After an initial setup process, Hitachi Dynamic Tiering monitors data access in real time. It makes decisions on moving data between the available storage tiers based on actual use. Using this approach, Hitachi Dynamic Tiering improves the availability and performance of your storage systems and the applications using that storage.

Hitachi Dynamic Tiering on Hitachi Virtual Storage Platform allows a single pool to contain tiers made up of differently-arranged RAID groups using any type of disk. It manages data migration between the various tiers within a pool automatically. This eliminates most user management of storage tiers within a storage system, and maintains peak performance under dynamic conditions without storage administrator intervention.

Hitachi Unified Compute Platform Director

Hitachi Unified Compute Platform Director offers the following:

- Full parity across the RESTful API
- Command line interface
- Graphical user interface

Director software on the Unified Compute Platform Pro integrates directly into VMware vSphere. It provides unified end-to-end infrastructure orchestration within a single interface.

Unified Compute Platform Director leverages your existing storage in one of two ways:

- Connect to your existing Hitachi Virtual Storage Platform or Hitachi Unified Storage VM
- Virtualize other storage arrays that you have from other vendors using Virtual Storage Platform, Virtual Storage Platform G1000 or Unified Storage VM

Unified Compute Platform Director provides the following benefits:

- Centralization and automation of compute, storage, and networking components
- Significant reduction of time to value and operational costs across data centers
- Faster deployment of converged infrastructure with more efficient resource allocation
- Provides a foundation for the journey to the software defined datacenter using full support of the RESTful API

Hitachi Compute Blade 500

[Hitachi Compute Blade 500](#) combines the high-end features with the high compute density and adaptable architecture you need to lower costs and protect investment. Safely mix a wide variety of application workloads on a highly reliable, scalable, and flexible platform. Add server management and system monitoring at no cost with Hitachi Compute Systems Manager, which can seamlessly integrate with Hitachi Command Suite in IT environments using Hitachi storage.

Hitachi Compute Blade 2000

[Hitachi Compute Blade 2000](#) is an enterprise-class blade server platform. It features the following:

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Embedded logical partition virtualization
- Configuration flexibility
- Sustainable power-saving capabilities
- Fast server failure recovery using a N+1 cold standby design that allows replacing failed servers within minutes

Hitachi embeds logical partitioning virtualization in the firmware of the Hitachi Compute Blade 2000 server blades. This proven, mainframe-class technology combines logical partitioning expertise from Hitachi with VT technologies from Intel to improve performance, reliability, and security. Embedded logical partition virtualization does not degrade application performance and does not require the purchase and installation of additional components.

Hitachi Compute Rack 210H

Hitachi Compute Rack 210H is midrange rack mountable server platform, providing advanced systems management and redundancy options. It is data center friendly, with a 1U footprint, while delivering the performance that is required to meet enterprise-level challenges.

The benefits of Hitachi Compute Rack 210H are the following:

- Web-based management interface
- RAID level configuration, with up to six 2.5 inch internal drives
- Sustainable power-saving capabilities
- Configuration flexibility to meet business needs
- Dense 1U rack mountable design
- 1U rack-mountable server

Hitachi Compute Rack 220H

Hitachi Compute Rack 220H is an Intel Xeon processor-based midrange rack mountable server platform, providing advanced systems management and redundancy options. It is data center friendly, with a 2U footprint, while delivering the performance that is required to meet enterprise level challenges.

The benefits of Hitachi Compute Rack 220H are the following:

- Web-based management interface
- RAID level configuration, with up to six 3.5 inch internal drives
- Sustainable power-saving capabilities
- 2 socket Intel based server
- Configuration flexibility to meet business needs
- Dense 2U rack mountable design
- 2U rack-mountable server

Hitachi Data Instance Manager

Hitachi Data Instance Manager (HDIM) is a unified, easy-to-use, policy and work flow based solution that centralizes multiple data protection capabilities for file, SQL® and Microsoft® Exchange® in Microsoft and Linux mid-tier environments, including remote offices. It includes backup, continuous data protection, archive, snapshot, de-duplication, data security and more, with top-down policy management. HDIM protects data and reduces, controls and manages copy data, or data instances. It replaces or augments legacy data protection point-solutions and eliminates backup windows, improves service levels, reduces backup data volumes, supports compliance requirements and improves application uptime, all while helping to reduce costs.

- Eliminate your backup window with a unified continuous data protection and archive solution
- Meet service levels and reduce costs using fewer resources to manage more data with a holistic approach to data protection.
- Real-time data capture enables maximum protection and consistent snapshots for instant recovery and application consistency

Hitachi Data Protection Suite, Powered by CommVault, with Hitachi Thin Image

Hitachi Data Protection Suite, powered by CommVault, is a tightly integrated blend of snapshot, replication and persistent copies that are application-aware, secure, de-duplicated, managed and accessible through a single, unified platform. This unified approach results in an easy, automated, and low cost data management solution, at a fraction of the time, effort, and money required by separate point products.

Today's dynamic storage environments are too complex to manage with native tools, loosely integrated software packages and disjointed processes. Hitachi Data Protection Suite delivers powerful data protection capabilities within a truly unified platform. The Data Protection Suite's automated, policy-driven IntelliSnap supports leading storage platforms and applications. It also helps users protect hundreds of virtual machines in minutes and easily scales to thousands of VMs and physical servers.

- Ensures 100% application consistent snapshots and shifts the data protection workload from the virtual platform to storage system
- Manages and protects hundreds of VMs within a single management interface with minimal impact on production environments
- Enables robust private and public cloud solutions
- Provides universal de-duplication—starting from the source all the way to tape, globally across sites—across VMs

Tightly integrates with Hitachi Content Platform to provide intelligent, policy-based data movement that enables seamless tiering to the cloud.

Hitachi Thin Image

An essential component of data backup and protection solutions is the ability to quickly and easily copy data. Hitachi Thin Image (HTI) snapshot provides logical, change-based, point-in-time data replication within Hitachi storage systems for immediate business use. Business usage can include data backup and rapid recovery operations, as well as decision support, information processing, and software testing and development.

- Maximum capacity of 5.0 PB enables larger data sets or more virtual machines to be protected
- Maximum snapshots increased to 1024 for greater snapshot frequency and/or longer retention periods
- Asynchronous operation greatly improves response time to host
- Enhanced for super-fast data recovery performance

HTI has the following benefits:

- Reduce recovery time from data corruption or human errors while maximizing Hitachi disk storage capacity
- Achieve frequent and nondisruptive data backup operations while critical applications run unaffected
- Accelerate application testing and deployment with always-available copies of current production information
- Significantly reduce or eliminate backup window time requirements
- Improve operational efficiency by allowing multiple processes to run in parallel with access to the same information

Hitachi Storage Plug-Ins for VMware vCenter

Hitachi Data Systems has storage plug-ins for both the VMware vSphere Client and the VMware vSphere Web Client.

Storage Manager for VMware vCenter 2.5

Storage Manager for VMware vCenter operates as a plug-in in the vSphere Client and displays the configuration information of the Hitachi storage systems (AMS2000 series, HUS100 series, USP V/VM, VSP, or HUS VM) attached to the vSphere environment.

Storage Manager for VMware vCenter provides the following functions in cooperation with the vCenter Server:

- The View function is the function used to display the storage system information registered in Storage Manager, the datastore on ESX Server using the storage system, and virtual machine information.
- The Provision Datastore function is the function to orchestrate in a single operation the steps required to create a datastore on an ESX Server using the LU of the storage system.
- Snapshot Datastore Function is the function used when a virtual machine using a template is copied. Manual operation can be eliminated by using the Snapshot Datastore function.
- Delete Datastore function is the function to orchestrate in a single operation all the steps required to delete a datastore and the underlying LU from the storage system registered in Storage Manager.

Hitachi Storage Plug-in for VMware vCenter 3.0

The Hitachi Storage Plug-in operates as a plug-in for the vSphere Web Client and displays the configuration information of the Hitachi storage systems attached to the vSphere environment.

Additionally, Hitachi Storage Plug-in provides the following functions in cooperation with the vCenter Server:

- The View function is the function used to display the storage system information registered in Storage Manager, the datastore on ESX Server using the storage system, and virtual machine information.
- The Provision Datastore function is the function to orchestrate in a single operation the steps required to create a datastore on an ESX Server using the LU of the storage system.
- Delete Datastore function is the function to orchestrate in a single operation all the steps required to delete a datastore and the underlying LU from the storage system registered in Storage Manager.

Hitachi Compute Plug-in for VMware vCenter

Hitachi Compute Plug-in for VMware vCenter is a vSphere Web Client plug-in that allows Hitachi Compute Blade 500 and 2000 server information to be integrated in to VMware vCenter.

VMware vSphere 5.5

[VMware vSphere 5.5](#) is a virtualization platform that provides a datacenter infrastructure. It features vSphere Distributed Resource Scheduler (DRS), high availability, and fault tolerance.

VMware vSphere 5 has the following components:

- **ESXi** — A hypervisor that loads directly on a physical server. It partitions one physical machine into many virtual machines that share hardware resources.
- **vCenter Server** — Management of the vSphere environment through a single user interface. With vCenter, there are features available such as vMotion, Storage vMotion, Storage Distributed Resource Scheduler, High Availability, and Fault Tolerance.

Hitachi Storage Provider for VMware vCenter, and VMware vSphere API for Storage Awareness

Hitachi Storage Provider for VMware vCenter is Hitachi Data Systems implementation of VMware vSphere API for Storage Awareness (VASA). Hitachi Storage Provider for VMware vCenter provides vSphere with information about Hitachi Data Systems storage arrays for tighter integration between Hitachi Data Systems storage and the virtual infrastructure. The shared information includes details on storage virtualization, including health status, configuration, capacity and thin provisioning. This level of detail can now be passed up through the VMware vCenter Server to the end user:

- Gather information around specific disk array features (such as snapshot, deduplication, replication state, RAID levels, and thin or thick provisioned) and status (health, troubleshooting and so on).
- Make it easy for administrators to select the most appropriate storage resources when a virtual machine is created based on high-level details around capacity, performance and capabilities of the underlying disk.
- Integrate with other vSphere features, like VMware vSphere Storage DRS™ and Profile-Driven Storage.

VMware vStorage APIs for Array Integration

VMware vStorage APIs for Array Integration (VAAI) allow vSphere environments to use advanced features of Hitachi storage arrays. Using vStorage APIs provides a way to use the advanced storage capabilities of Hitachi Unified Storage from within the VMware interface. Processing is on the storage infrastructure directly.

These performance enhancements move the I/O load from the dependent vCenter host platform into the storage controller. Instead of slowing processing by consuming processing power, memory, and bandwidth, moving these operations speed operations, shift potential bottlenecks of resource constraint, and free virtualization management for more critical tasks.

When used with vSphere 5.x, Hitachi Unified Storage supports the following API primitives:

- Full copy—This primitive enables the storage system to make full copies of data within the storage system without having the ESX host read and write the data.
- Block zeroing—This primitive enables storage systems to zero out a large number of blocks to speed provisioning of virtual machines.
- Hardware-assisted locking—This primitive provides an alternative means to protect the metadata for VMFS cluster file systems, thereby improving the scalability of large ESX host farms sharing a datastore.
- Thin provisioning stun—This primitive enables the storage system to notify the ESX host when thin provisioned volumes reach certain capacity utilization threshold. When enabled, this allows the ESX host to take preventive measures to maintain virtual machine integrity.
- UNMAP - Enables an ESXi host to inform the Hitachi Data Systems storage array that space can be reclaimed that previously had been occupied by a virtual machine that has been migrated to another datastore or deleted.

Performance and Scalability Best Practices

The following sections describe best practices for optimizing your storage infrastructure to meet your performance and scalability, requirements.

VMware vSphere Storage APIs for Array Integration

To use VAAI with a Hitachi Data Systems storage array, the following are required:

- ESX or ESXi 4.1 or later (ESXi 5 for NAS support)
- ESX or ESXi Enterprise or Enterprise Plus licensing
- A Hitachi Data Systems storage array with a microcode that supports VAAI (Check the microcode documentation or with a Hitachi Data Systems representative).

To enable VAAI on the Hitachi Data Systems VSP, VSP G1000, and HUS VM choose Host Mode Options 54 and 63. To Enable VAAI on HUS, the Host Group Platform setting should be set to VMware.

To verify that VAAI is enabled on your ESX or ESXi 4.x hosts from the service console.

```
esxcfg-advcfg -g /DataMover/HardwareAcceleratedMove
esxcfg-advcfg -g /DataMover/HardwareAcceleratedInt
esxcfg-advcfg -g /VMFS3/HardwareAcceleratedLocking
```

To verify that VAAI is enabled on your ESX or ESXi 5.x hosts from the service console.

```
esxcli system settings advanced list -o \ /DataMover/HardwareAcceleratedMove
esxcli system settings advanced list -o \ /DataMover/HardwareAcceleratedInt
esxcli system settings advanced list -o \ /VMFS3/HardwareAcceleratedLocking
```

If VAAI is enabled on the ESX or ESXi hosts the value for Int Value: will be set to 1.

To verify that VAAI is enabled on your ESX or ESXi hosts from the vSphere Client:

- Select the ESX or ESXi host in vCenter
- With the ESX or ESXi host selected click the **Configuration** tab
- On the Configuration tab, under software, click **Advanced Settings**

The following options should have their values set to 1 for VAAI to be enabled:

```
DataMover.HardwareAcceleratedMove
DataMover.HardwareAcceleratedInt
VMFS3.HardwareAcceleratedLocking
```

For more information on VAAI on Hitachi Data Systems storage platforms see [the Lab Validation Report about VAAI on the VSP](#) and [Lab Validation Report about VAAI on the HUS 150](#).

VMware vSphere Storage I/O Control

The best practice from Hitachi Data Systems is to set the default HBA LUN queue depth to 32.

Setting the HBA LUN queue depth greater than 32 may overload the storage processors on your storage systems. This depends on the load from the ESX or ESXi hosts.

In a mixed environment, with ESXi hosts as well as bare metal Microsoft Windows and/or Linux hosts, setting the HBA LUN queue depth for the ESXi hosts greater than 32 can negatively impact the performance of the bare metal hosts.

Performance improvements were archived with Hitachi block storage solutions and Storage I/O Control (SIOC) when the default HBA LUN queue depth was set to 64 or higher. The best results were found at 128.

Table 3. HBA LUN Queue Depth with Qlogic for ESX and ESXi Versions 3.5 through 5.5

ESX or ESXi Version	HBA LUN Queue Depth
3.5 with driver version 6.04	16
3.5 with driver version 6.07 and 7.x	32
4.0 through 4.1	32
5.0 through 5.5	64

The default HBA LUN queue depth for Emulex and Brocade HBA adapters is 32. Only 30 are available for I/O data on the Emulex HBA.

If your ESX and ESXi hosts are using a mixture of HBA vendors, the recommendation is to use a consistent HBA queue depth across all the hosts in a cluster.

When adjusting the HBA LUN queue depth, do not start with the highest queue depth value. It is true that this may give SIOC the greatest room to equally distribute I/O across the virtual machines on a datastore. However, increasing the HBA LUN queue depth too high may overload the storage processors on your storage systems, depending on the load from the ESX or ESXi hosts.

The recommendation is to start with a queue depth of 32. Monitor the performance of the virtual machines and the storage systems. If needed, increase the HBA LUN queue depth to 64 and monitor again. Repeat this change-and-monitor process each time you increase the HBA LUN queue depth.

Note — Increasing the HBA LUN queue depth beyond 32 can have a negative effect on the performance of the storage system and other bare metal systems running other operating systems.

For more information on Storage I/O Control on Hitachi Data Systems storage platforms see the following:

- [Advantages of using VMware Storage I/O Control with UCP Select for VMware vSphere using Hitachi Unified Storage VM Lab Validation Report](#)
- [Advantages of using VMware Storage I/O Control with Hitachi Unified Compute Platform Select with VMware vSphere using Hitachi Unified Storage 150 Lab Validation Report](#)

To change the default HBA LUN queue depth on your ESX and ESXi hosts, follow the instructions in the following: VMware ESXi 5.5 Hosts:

- “VMware ESXi 5.5 Hosts” on page 14
- “VMware ESXi 5.0 and VMware ESX 5.1 Hosts” on page 15

VMware ESXi 5.5 Hosts

To change the default HBA LUN queue depth on a VMware ESXi 5.5 host, do the following.

1. Verify the currently loaded HBA module by typing one of these commands on the service console:

- **QLogic**

```
esxcli system module list | grep qln
```

- **Emulex**

```
esxcli system module list | grep lpfc
```

- **Brocade**

```
esxcli system module list | grep bfa
```

- **Hitachi**

```
esxcli system module list | grep hfc
```

2. Change the HBA LUN queue depth.

Note — The examples show the QLogic qla2xxx and Emulex lpfc820 modules. Use the command for the HBA module that you verified.

- **QLogic**

```
esxcli system module parameters set -p ql2xmaxqdepth=64 -m qlnativefc
```

- **Emulex**

```
esxcli system module parameters set -p lpfc0_lun_queue_depth=64 -m lpfc
```

- **Brocade**

```
esxcli system module parameters set -p bfa_lun_queue_depth=64 -m bfa
```

- **Hitachi**

```
esxcli system module parameters set -p hfc_lun_queue_depth=64 -m hfc
```

3. Reboot your host.

4. To confirm that your changes have been applied, run this command:

```
esxcli system module parameters list -m [driver]
```

Where [driver] is your QLogic, Emulex, Brocade, or Hitachi adapter driver module, such as lpfc820, qla2xxx, hfc, or bfa.

VMware ESXi 5.0 and VMware ESX 5.1 Hosts

To change the default HBA LUN queue depth on a VMware ESXi 5.5 host, do the following.

1. Verify the currently loaded HBA module by typing one of these commands on the service console:
 - **QLogic**
`esxcli system module list | grep qla`
 - **Emulex**
`esxcli system module list | grep lpfc`
 - **Brocade**
`esxcli system module list | grep bfa`
 - **Hitachi**
`esxcli system module list | grep hfc`
2. Change the HBA LUN queue depth.

Note — The examples show the QLogic qla2xxx and Emulex lpfc820 modules. Use the command for the HBA module that you verified.

- **QLogic**
`esxcli system module parameters set -p ql2xmaxqdepth=64 -m qla2xxx`
 - **Emulex**
`esxcli system module parameters set -p lpfc0_lun_queue_depth=64 -m lpfc820`
 - **Brocade**
`esxcli system module parameters set -p bfa_lun_queue_depth=64 -m bfa`
 - **Hitachi**
`esxcli system module parameters set -p hfc_lun_queue_depth=64 -m hfc`
3. Reboot your host.
 4. To confirm that your changes have been applied, run this command:
`esxcli system module parameters list -m [driver]`

Where [driver] is your QLogic, Emulex, Brocade, or Hitachi adapter driver module, such as lpfc820, qla2xxx, hfc, or bfa.

VMware vSphere Storage DRS

VMware vSphere Storage DRS continuously balances storage space usage and storage I/O load while avoiding resource bottlenecks to meet application service levels. Using Storage DRS allows you to:

- Take advantage of new capacity as it is added to a storage cluster
- Improve service levels for applications running on virtual machines whose storage is in a Storage DRS enabled storage cluster.
- Decrease VMware vSphere administrator overhead by automating storage capacity and load balancing.

As a best practice, do not include datastores that have hardware acceleration (VAAI) enabled in the same datastore cluster as datastores that do not have hardware acceleration enabled. Datastores in a datastore cluster must be homogeneous to guarantee hardware acceleration-supported behavior.

All hosts attached to the datastore must be ESXi 5.0 and later.

If this SDRS deployment will be new for the organization or the workload, it is recommended to first set the automation level to Manual Mode. Then manually apply the recommendations and monitor the impact on the application workloads. When the administrator is comfortable with the SDRS recommendations the storage cluster can be changed to Fully Automatic. The aggressiveness of SDRS may need to be adjusted higher or lower.

Disk Alignment

Partition alignment is the proper alignment of the virtual machine (VM) partitions to the boundaries of the storage array (such as a hard disk, flash, or RAID volume). Proper partition alignment ensures the best performance during data access. Incorrect partition alignment will cause degraded performance by causing a single I/O request from the VM OS into 2 or more I/O requests on the underlying storage system.

To avoid partition misalignment and ensure proper partition alignment the following recommendations are made.

- When working with vSphere 5.0 or greater and an ESXi 5.0 or greater, use the vSphere Client to create VMFS partitions, since ESXi 5.0 and above automatically aligns VMFS3 or VMFS5 partitions along the 1MB boundary.
- If a VMFS3 partition was created using an earlier version of ESX/ESXi than 5.0 that aligned along the 64KB boundary, and that file system is then upgraded to VMFS5, it will retain its 64KB alignment. 1MB alignment can be obtained by deleting the partition and recreating it using the vSphere Client and an ESXi 5.0 or later host.
- To manually align your VMFS partitions use a starting block address that is a multiple of 8KB.
- To correct an existing VM that is misaligned VMware Converter can be used to run the existing VM through the converter and the new resulting VM will be properly aligned.
 - In VMware Converter from the *Data copy type* drop-down menu, select *Select volumes to copy*.
 - Select the *Create optimized partition layout* check box.

VMDK Provisioning Type

When creating a virtual disk, migrating a virtual machine, or cloning a virtual machine to a template you can specify a provisioning type for the virtual disk file. You can use Storage vMotion to convert virtual disks from one provisioning type to another.

VMFS datastores support the following disk provisioning policies.

Thick Provision Lazy Zero

This is the default virtual disk provisioning type. Space required for the virtual disk is allocated when the disk is created. Data remaining on the underlying storage system is not erased during creation, but is zeroed out on first write from the virtual machine. Virtual machines do not read stale data from the physical device. When using Hitachi Dynamic Provisioning (HDP), the un-written to blocks are not allocated to the virtual disk until they are written to the first time by the virtual machines (VMs) guest OS. This allows space savings on the storage system and saves the VMkernel overhead of having to grow the virtual disk file.

Thick Provision Eager Zero

A type of virtual disk supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. Unlike the thick provision lazy zeroed format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take longer to create virtual disks in this format than to create other types of disks. Also when using HDP all the space for the virtual disk is allocated at the time of creation. There is now cost savings for the un-written to part of the virtual disk. A 40 GB virtual disk takes up 40 GB of space on the storage system, even before a guest OS is installed or data is written.

Thin Provision

Use this type of virtual disk to save storage space. With the thin disk, you provision as much datastore space as the disk would require based on the value that you enter for the virtual disk size. However, the thin disk starts small and at first, uses only as much datastore space as the disk needs for its initial operations. If the thin disk needs more space later, it can grow to its maximum capacity and occupy the entire datastore space provisioned to it.

When using HDP is it recommended to use Thick Provisioned Lazy Zeroed virtual disks. This will allow the space and cost savings of Thin Provisioned virtual disks but use fewer resources on the ESX / ESXi hosts, as the host do not take a CPU hit when trying to zero blocks on first write or grow the underlying thin provisioned VMDK file of the virtual disk.

For more information about using Hitachi Dynamic Provisioning (HDP) with VMware vSphere see [Using VMware vSphere 5 with Hitachi Dynamic Provisioning on Hitachi Unified Storage](#).

Hitachi Accelerated Flash

As the storage I/O in most virtualized environments are heavily random it is recommended to configure Hitachi Accelerated Flash module as RAID-10 (2D+2D) for maximum performance.

Hitachi Accelerated Flash is ideal for I/O intensive applications that are also sensitive to latency or where very low latency is desired.

Hitachi Accelerated Flash is also an ideal solution in environments where a minimal foot print is required or preferred in the datacenter. In testing in Hitachi's lab, 216 SAS drives in nine 2U disk trays were replaced by four Hitachi Accelerated Flash modules that took only a third of a tray during testing.

Hitachi Accelerated Flash simplifies the storage pool and LUN management, as all mixed workloads can be placed together into a fewer number pools.

When using Hitachi Accelerated Flash with Hitachi Dynamic Provisioning in a VMware vSphere environment, following settings are recommended:

- Flash Acceleration Microcode - Enabled
- SOM 917 on Hitachi Unified Storage VM – Enable
 - SOM 917 balances page placement across parity groups rather than pool volumes.
- DPVOL (LU) Ownership – Distribute evenly to all MPUs
- Host Group Mode – 21[VMware Extension], and enable the following:
 - 54-Support Option for the EXTENDED COPY command
 - 63-Support option for vStorage APIs based on T10 standards
- Multipathing Policy in ESXi Host – Round Robin
- Virtual Machine Placement on each ESXi host – Place evenly on the LUs to balance the I/O loads

For more information on Hitachi Accelerated Flash see [Hitachi Unified Compute Platform Select for VMware vSphere using Hitachi Accelerated Flash on Hitachi Unified Storage VM](#).

Reliability and Scalability Best Practices

SAN design should begin with reliability and scalability in mind. A well designed SAN will be able to recover quickly in the event of a single device failure. Also a well-designed SAN will easily grow as the demands of the infrastructure that it serves increases.

Multipathing

Multipathing allows an ESX or ESXi host to use more than one physical path between the ESX or ESXi host and the storage array.

Multipathing provides load balancing, which is the process of distributing I/O across multiple physical paths, reduces or removing potential bottlenecks.

Multipathing also provides redundancy and fault tolerance in case of a failure of any element in the SAN network such as an adapter, switch, or cable. The ESX or ESXi host can switch to another physical path, which does not use the failed component. This process of path switching to avoid failed components is known as path failover.

To support path switching with an FC SAN, the ESXi host typically has two or more HBAs available from which the storage array can be reached and for full fault tolerance uses two or more switches. Additionally for full fault tolerance two storage processors on the Hitachi storage arrays should be utilized so that the HBA can use a different path to reach the disk array.

Path Selection / Multipathing policies supported by the ESX and ESXi hosts are:

- **Most Recently Used (VMware)** - Uses the last successfully used path. If the last successfully used path is not available then path failover occurs and the next available path is used. This new path continues to be used until it is no longer available, even if the previously used path becomes available again.
- **Round Robin (VMware)** - Sends a set number of I/O down the first available path, then sends the same set number of I/O down the next available path. This repeats through all available paths then start over again and repeats. If a path becomes unavailable it is skipped over until it becomes available again.
- **Fixed (VMware)** - Has a preferred path that can be set by the vCenter administrator. This path is used until it becomes unavailable. Then it fails over to the next available path until it becomes unavailable, in which case it fails over to the next available path, or until the preferred path becomes available again. If the preferred path does become available again, then the system fails back to the preferred path.
- **Hitachi Dynamic Link Manager** - VMware ESX and ESXi also support third party Path Selection policies. Hitachi Dynamic Link Manager (HDLM) is Hitachi's multipathing software that integrates with Hitachi High Availability Manager to provide load balancing and path failover capabilities for servers. Consider this when implementing a VMware Metro Storage Cluster environment.

The recommended multipathing policy is Round Robin (VMware). This multipathing policy takes advantage of all available paths and bandwidth. By taking advantage of all available paths, maximum performance from the SAN infrastructure is assured.

VMware ESXi Round Robin Path Selection Plug-in Path Rotation IOPS Limit

ESXi Round Robin Path Selection Plug-in (PSP) balances the load across all active storage paths. A path is selected and used until a specific quantity of I/O has been transferred.

The I/O quantity at which a path change is triggered is known as the limit. After reaching that I/O limit, the PSP selects the next path in the list.

The default I/O limit is 1000, but can be adjusted if needed to improve performance. Specifically, it can be adjusted to reduce latency seen by the ESXi host when the storage system is not seeing latency.

The recommended PSP limit is 1000. Setting the limit to 1 has not shown any improvement in IOPS or latency when testing with Hitachi storage systems.

Redundant HBA

In a VMware vSphere environment the ESX or ESXi hosts should have 2 or more HBAs. Not only will this allow for greater I/O throughput on the SAN as more paths are available when using the Round Robin (VMware) multipathing policy, multiple HBAs also allow for redundancy and greater reliability in case of a component failure.

Redundant Fibre Channel Switches

When designing and building a reliable and scalable SAN environment, multiple Fibre Channel switches are recommended. Each ESX or ESXi host (the port in the Fibre Channel HBA is referred to as the initiator) should have an equal number of connections to each Fibre Channel switch and each Hitachi storage array (the port in the Hitachi storage arrays storage processor is referred to as the target) should have an equal number of connections from each storage processor to each switch.

This SAN Fibre Channel switch configuration ensures that a single switch failure will not leave an ESX or ESXi host unable to connect to a datastore and therefore unable to continue running the virtual machines on those datastores.

It is recommended that the Fibre Channel switches not be up-linked to each other creating separate Fibre Channel networks. This insures that conditions on the Fibre Channel network, for example a mal-functioning HBA, do not affect traffic on the other Fibre Channel network. This helps ensure great reliability.

Zoning

Zoning is used to enable access control in a SAN environment. Through zoning a SAN administrator can configure which HBA WWPNs (World Wide Port Names) on the ESX or ESXi host can connect to which WWPNs on the Hitachi storage array's storage processors.

Zoning Configurations

Zoning can be broken down into three different configurations. There are Single Initiator to Single Target zones, Single Initiator to Multiple Target Zones and Multi Initiator Zones.

- **Single Initiator to Single Target zones** - Allows one initiator to be zoned to one target. This configuration is the most restrictive and requires the most administrative overhead, but does give us the most resilient configuration as traffic originating from another Initiator on the SAN will have less impact than the initiator in this zone.
- **Single Initiator to Multiple Target Zones** - Allows one initiator to be zoned to multiple targets in a single zone.
- **Multi Initiator Zones** - Allows multiple initiators to be zoned to multiple targets in a single zone. This exposes all initiators and targets to all traffic in the zone. Events such as a malfunctioning HBA could affect all initiators and targets in the zone and either negatively effect performance for all or bring down the Fibre Channel network completely. This configuration is never recommended.

A Single Initiator to Single Target zone is recommended as it is the most secure and stable, but Single Initiator to Multiple Target zone is also acceptable for ease of administration.

Host Group Configuration

To grant a host access to an LDEV, assign a logical unit number (LUN) within a host group. This section describes the settings and LUN mapping for host group configurations.

Fibre Channel Port Options

If a Fibre Channel port is connected via a SAN switch or director, you must change the following settings in Hitachi Storage Navigator software:

- Port security — Set to Enable; allows multiple host groups on the Fibre Channel port.
- Fibre Channel switch — Set to Enable; allows connection to a Fibre Channel switch or director.
- Connection type — Set to P-to-P; allows a point-to-point connection to a Fibre Channel switch or director.

Hitachi Data Systems recommends that any configuration applied to a port in cluster 1 also is applied to the port in cluster 2 in the same location. For example, if you create a host group for a host on port CL1-A, also create a host group for that host on port CL2-A.

One Host Group per ESX Host, Standalone Host Configuration

If you plan to deploy ESX hosts in a standalone configuration, each host's WWNs can be in its own host group. This approach provides granular control over LUN presentation to ESX hosts. This is the best practice for SAN boot environments, because ESX hosts do not have access to other ESX hosts' boot LUNs. However, in a cluster environment this approach can be an administrative challenge because keeping track of which ESX hosts are in a cluster can be difficult. When multiple ESX hosts need to access the same LDEV for clustering purposes, the LDEV must be added to each host group. This operation is error prone and might lead to confusion. If you have numerous ESX hosts, this approach can be tedious.

One Host Group per Cluster, Cluster Host Configuration

vSphere features such as vMotion, DRS, HA, FT and Storage vMotion require shared storage across the ESX hosts. Many of these features require that the same LUNs are presented to all ESX hosts participating in these cluster functions. If you plan to use ESX hosts with these features, create host groups with clustering in mind. Place all of the WWNs for the clustered ESX hosts in a single host group. This ensures that when LDEVs are added to the host group, all ESX hosts see the same LUNs. This creates consistency with LUN presentation across all hosts.

Host Group Options

On the Virtual Storage Platform, host groups are created using Hitachi Storage Navigator software. The host mode is set for the host operating system. Select the host mode of 21[VMware] for all host groups to be used by ESX or ESXi hosts.

Usability Best Practices

Storage Manager for vCenter

Storage Manager for VMware vCenter is a VMware vCenter 4.1 and 5.0 plug-in that provides integrated management of Hitachi storage subsystems within vCenter. The plug-in allows the vCenter administrator to automatically configure ESX hosts to Hitachi storage subsystems.

The administrator can provision, map, and delete storage logical drives from the Hitachi storage subsystem to the ESX host, and view the vCenter datastores to Hitachi storage subsystem logical drives.

Storage Manager provides the following functions in cooperation with the vCenter Server.

- View Function
- Provision Datastore Function
- Snapshot Datastore Function
- Delete Datastore Function

View Function

The View function is the function used to display the storage system information registered in Storage Manager, the datastore on ESX Server using the storage system, and virtual machine information.

The View function facilitates the confirmation of the following:

- The ESX Server connected to the storage system, and an LU that is to be used with ESX Server
- The storage system used with the datastore and virtual machines

Provision Datastore Function

The Provision Datastore function is the function to orchestrate in a single operation, the steps required to create a datastore on an ESX Server using the LU of the storage system.

- (1) Create an LU in the storage system.
- (2) Assign the LU created to the Host Group or iSCSI Target.
- (3) Rescan the storage system with the storage adapter (HBA of ESX Server) to detect the LU created.
- (4) Create a datastore using the detected LU.

If the LU being provisioned is to be used as RDM, without creating a datastore, then only the procedures from (1) to (3) shown above are executed. The assignment of the RDM to the virtual machine needs to be performed manually.

Snapshot Datastore Function

When a virtual machine using a template is copied, manual operation can be eliminated by using the Snapshot Datastore function.

The Snapshot Datastore function is the function used to perform the following steps for multiple copying of a datastore including the template created in a P-VOL to a V-VOL.

- (1) Use the LUN where the copy source datastore is located as the P-VOL to create a V-VOL on the storage system.
- (2) Assign the created V-VOL to the specified Host Group or iSCSI target.
- (3) Rescan the storage array using the storage adapter (the HBA of ESX Server) to detect the created V-VOL.
- (4) Register the datastore of the detected V-VOL in the inventory.

The Snapshot Datastore function uses the Copy-on-Write SnapShot function of the HUS 100 series. Originally, the purpose of Copy-on-Write SnapShot function was to backup data. However, the Snapshot datastore function itself is a point in time reference that can be used for rolling back changes and as such cannot be used as a backup function. Also, the Snapshot Datastore function only supports HUS 100 series, and other storage systems are not supported.

Delete Datastore Function

Delete Datastore function is the function to orchestrate in a single operation all the steps required to delete a datastore and the underlying LU from the storage system registered in Storage Manager.

To delete a datastore, the following operations are performed:

- (1) Delete Datastore.
- (2) Unpresent the assignments to host group or iSCSI target of the LU used in the datastore.
- (3) Delete the LU used in the datastore.

Also, if you only delete the LU, perform the following procedures:

- (1) Unpresent the assignments to host group or iSCSI target of the LU.
- (2) Delete the LU.

The purpose of Delete Datastore function is to delete a datastore and its underlying LU if they were both created and provisioned via Storage Manager. Datastores (and their underlying LUs) which were not created using Storage Manager cannot be deleted.

Hitachi Compute Plug-in for VMware vCenter

Hitachi Compute Plug-in for VMware vCenter is a vSphere Web Client plug-in that allows Hitachi Compute Blade 500 and 2000 server information to be integrated in to VMware vCenter. Hitachi Compute Plug-in provides the following benefits :

- Detailed CB 500/2000 hardware information is displayed in the vSphere Web Client
- Information on the ESXi instances and VMs running on Hitachi CB 500/2000 servers is displayed in the vSphere Web Client
- The ability to perform power, LED, and remote console operations on CB 500/2000 servers is made available in the vSphere Web Client

Once the Hitachi Compute Plug-in is installed, a new option is available, as shown in Figure 1.

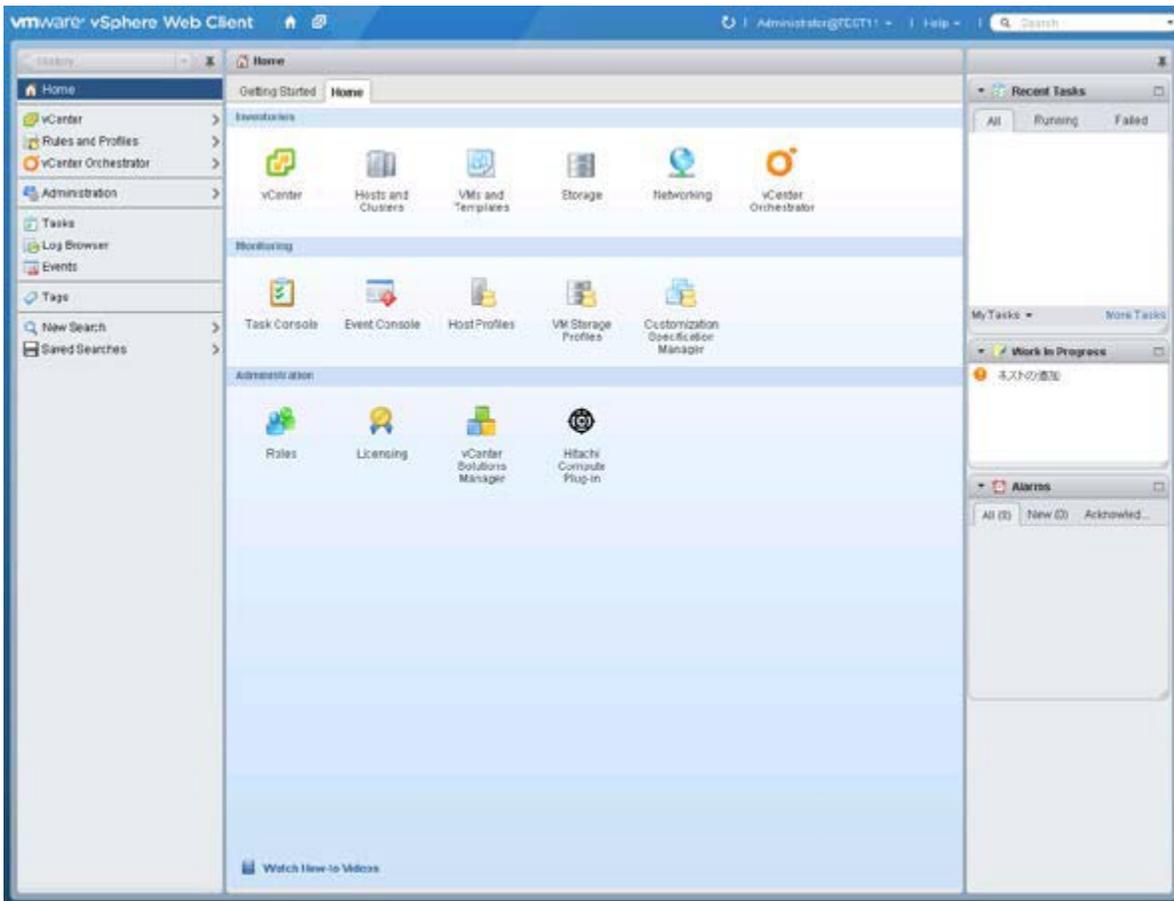


Figure 1

Once the Hitachi Compute Plug-In is selected, then the chassis or individual servers can be selected. When viewing the chassis you can see a summary of the chassis as shown in figure 2. If the Related Objects tab is selected then the server blades in the chassis, or the VMs running on that chassis, are visible as shown in Figure 2.

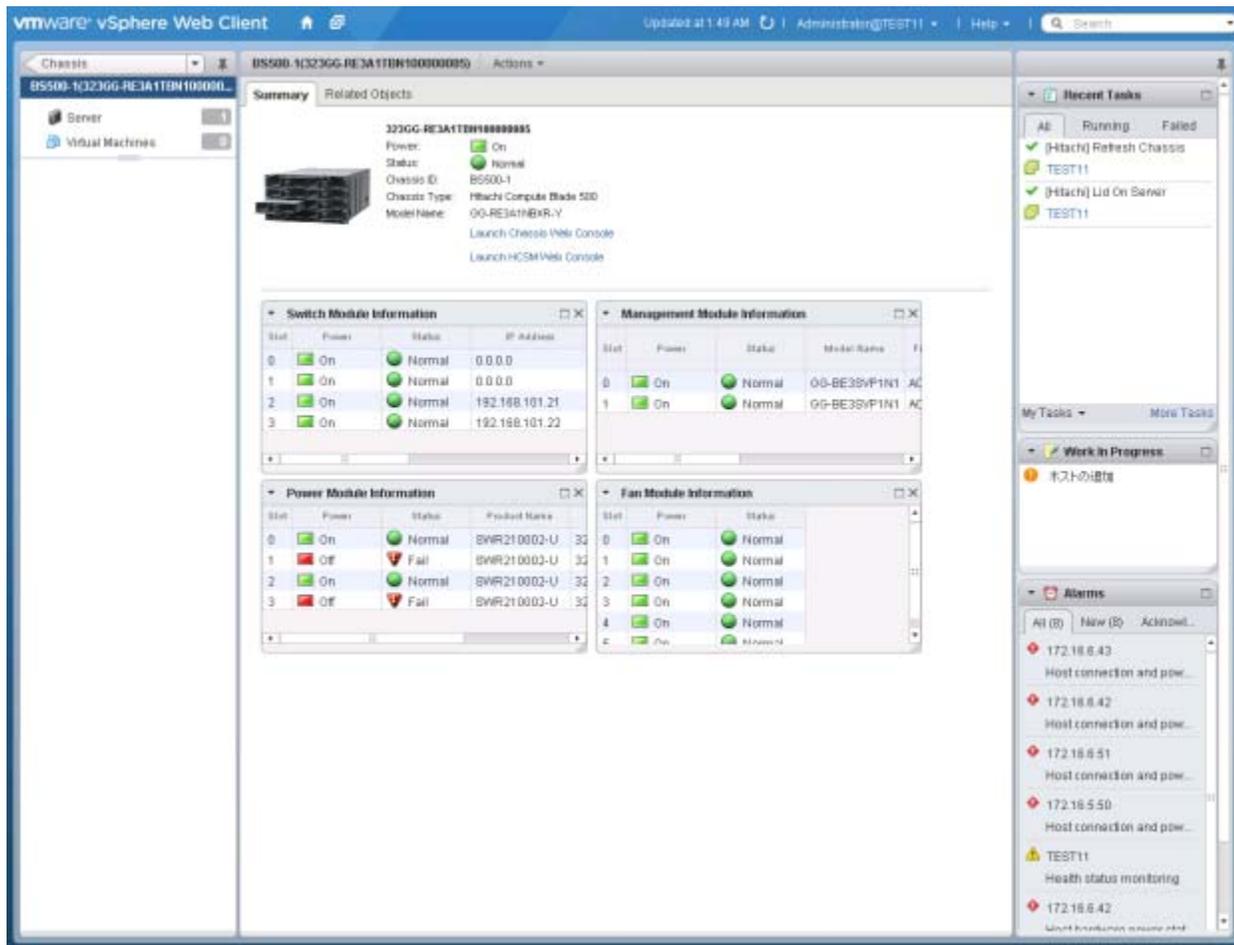


Figure 2

VMware vStorage API for Storage Awareness

Storage Provider for VMware vCenter is a VMware vStorage API for Storage Awareness, (VASA) Adapter that provides a standalone web service for VMware vCenter 5.1. This plug-in provides visibility from VMware vCenter into the Storage Array, and it enables VMware vCenter to be aware of physical storage resources and information based on topology, capability, and component state. This information can be used by vCenter Administrators to make more aware decisions or to create more effective policies toward Storage DRS and Profile Based Storage. Storage Manager for VMware vCenter is required to configure Storage Provider for VMware vCenter.

About vSphere API for Storage Awareness APIs Providers

VASA enables unprecedented coordination between VMware vCenter and storage arrays. It provides built-in storage insight to support intelligent VM storage provisioning, bolster storage troubleshooting and enable new DRS-related use cases for storage.

Hitachi Data Systems supports vSphere Storage APIs for Storage Awareness via the Storage Manager plug-in and the Storage Provider. The Storage Provider is compatible with VMware vSphere 5.x and supports storage array makes and models as described on the [VMware Compatibility Guide](#) and those listed below.

New in Version 02.1.0

- Support for vCenter v5.0 and vCenter v5.1 without installing Hitachi Storage Manager Plugin
- Configuration CLI via powershell for provider registration
- Addition of CAPABILITY profile for HNAS
- Bug fixes and architectural improvements

Figure 3 shows the information Hitachi Storage Provider for VMware vCenter provides to vCenter.

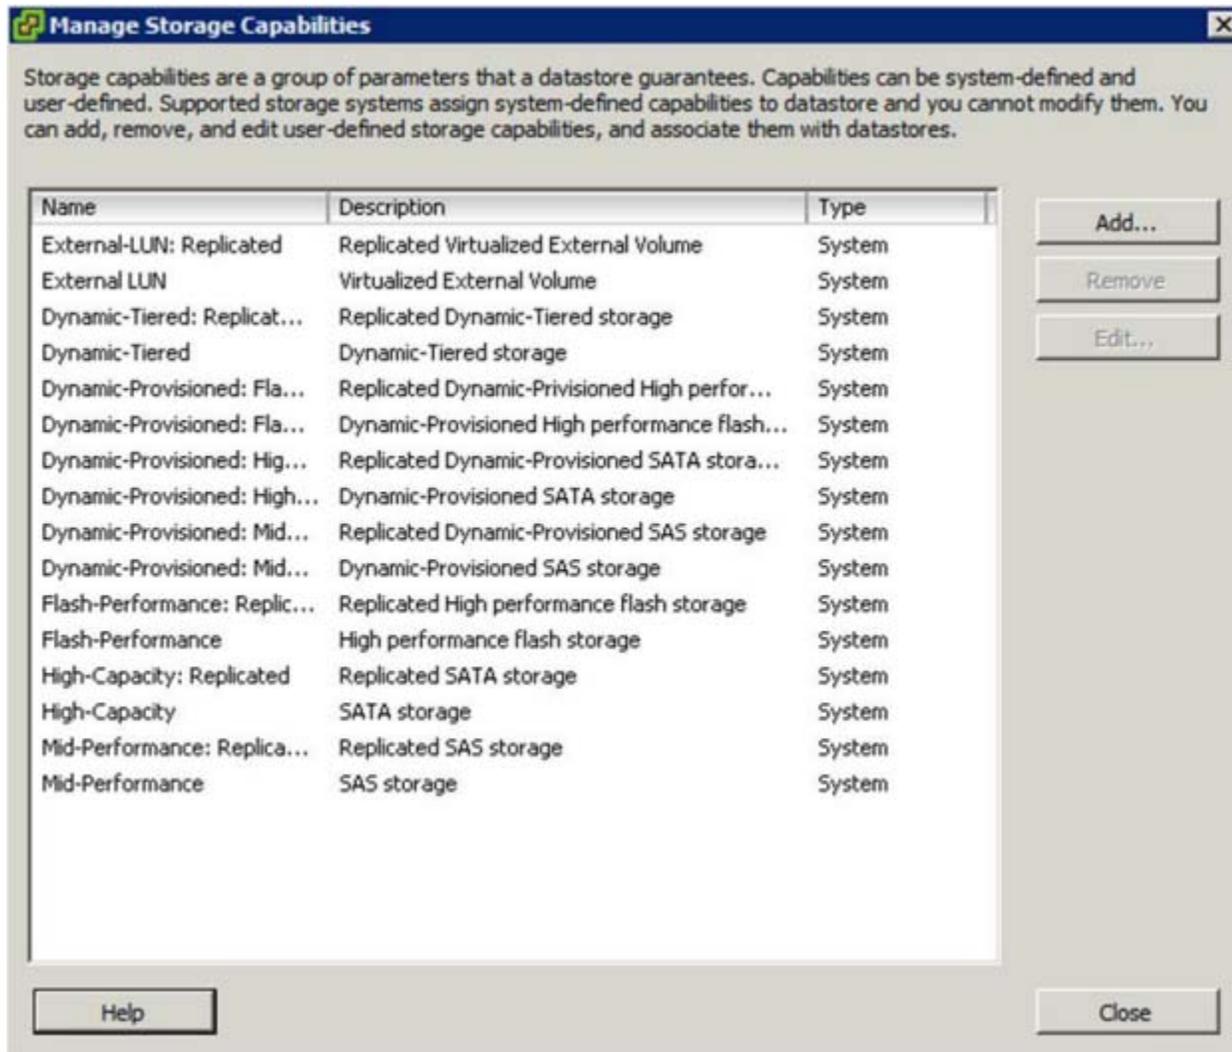


Figure 3

When the Storage Views tab is selected you can see the additional information passed to vCenter by the Hitachi Storage Provider for VMware vCenter, as shown in Figure 4.

System Capability	Capacity	Committed	Thin Provisioning	Storage Array	Identifier on Arr...	Volume Name
Mid-Performance	130.00 GB	130.00 GB	false	ASE-45.33 - VSP	52	VASAVSP
Mid-Performance	350.00 GB	350.00 GB	false	ASE-45.33 - VSP	51	HITACHI Fibre Chanr
External LUN	225.00 GB	225.00 GB	false	ASE-45.33 - VSP	36	HITACHI Fibre Chanr
Flash-Performance	225.00 GB	225.00 GB	false	ASE-45.33 - VSP	15	HITACHI Fibre Chanr
	100.00 GB	840.00 MB	true	ASE-45.33 - VSP	58	HITACHI Fibre Chanr
Dynamic-Provisioned: Mid-Performance	100.00 GB	840.00 MB	true	ASE-45.33 - VSP	57	Test Rename
Dynamic-Provisioned: Mid-Performance	100.00 GB	840.00 MB	true	ASE-45.33 - VSP	56	This is LUN 38
Mid-Performance	50.00 GB	50.00 GB	false	ASE-45.33 - VSP	40	HITACHI Fibre Chanr
Dynamic-Tiered	100.00 GB	840.00 MB	true	ASE-45.33 - VSP	55	HITACHI Fibre Chanr
Mid-Performance	75.00 GB	75.00 GB	false	ASE-45.33 - VSP	13	
	100.00 GB					HITACHI Fibre Chanr
	50.00 GB					HITACHI Fibre Chanr
	100.00 GB					HITACHI Fibre Chanr
	75.00 GB					HITACHI Fibre Chanr
	75.00 GB					HITACHI Fibre Chanr
	75.00 GB					HITACHI Fibre Chanr
	75.00 GB					SAPLVMOS_BOOT
	136.13 GB					Local DELL Disk (nas.
	100.00 GB					
	1.00 GB					
	800.00 GB					
	100.00 GB					
	50.00 GB					
	100.00 GB					

Figure 4

Hitachi Dynamic Tiering

Hitachi Dynamic Tiering (HDT) eliminates manual data classification and movement between storage tiers. This optimizes tiered storage usage while improving performance.

Instead of manually provisioning space from several storage technologies with different performance and cost characteristics, Hitachi Dynamic Tiering enables the management of multiple storage tiers as a single entity. By leveraging the thin provisioning and wide striping features of Hitachi Dynamic Provisioning, Hitachi Dynamic Tiering presents a virtual volume with embedded smart tiering. It monitors access and moves data at the 42 MB page level.

Breaking the volume into pages, Hitachi Dynamic Tiering automatically moves infrequently referenced pages to lower cost tiers of storage. Moving pages instead of entire data sets or files reduces the time and storage space required to migrate data.

Use Cases for Dynamic Tiering on vSphere

Virtual machines with very different workloads are often combined in a single DRS cluster in VMware environments. Segmenting these virtual machines on storage tiers is a labor-intensive administrative task. Dynamic Tiering can save you time and money by automatically placing less frequently accessed pages on lower-cost storage, and can improve performance by placing more frequently accessed pages on high-speed storage.

Hitachi Accelerated Flash with Hitachi Dynamic Tiering

With the wide spread adoption of flash and the proven performance of Hitachi Accelerated Flash, it is becoming more common that the top tier in an HDT volume will be Hitachi Accelerated Flash. Here are a few recommendations when using HDT with Hitachi Accelerated Flash:

- Flash Acceleration Microcode – Enable
- SOM 901 on Hitachi Virtual Storage Platform – Enable
 - SOM 901 permits an SSD/FMD tier to reach 100% performance utilization. By default, Hitachi Dynamic Tiering tries to keep all tiers at or below 60% performance utilization to improve response time. A SSD/FMD tier can deliver relatively good response times, even when 100% busy.
- SOM 917 on Hitachi Virtual Storage Platform – Enable
 - SOM 917 balances page placement across parity groups rather than pool volumes.
- DPVOL (LU) Ownership – Distribute evenly to all MPUs
- Host Group Mode – 21[VMware Extension], and enable following:
 - 54-Support Option for the EXTENDED COPY command
 - 63-Support option for vStorage APIs based on T10 standards
- Multipathing Policy in ESXi Host – Round Robin.
- Virtual Machine Placement on each ESXi host – Place evenly on the LUs to balance the I/O loads

For more information on Hitachi Dynamic Tiering with VMware vSphere [see Hitachi Unified Compute Platform Select for VMware vSphere using Hitachi Dynamic Tiering with Hitachi Accelerated Flash on Hitachi Virtual Storage Platform.](#)

Virtual Desktop Infrastructure Environment Best Practices

In a Virtual Desktop Infrastructure (VDI) environment, it is common to see the use of linked clones to conserve allocated disk space on your storage array. When the parent virtual machine is configured properly before the deployment of the desktops, this can be a great space saving technology.

Persistent desktops are becoming more prevalent in VDI environments within the corporate datacenter. Persistence provides the following benefits:

- Allows users to install custom applications
- Keeps user data locally on the desktop
- Guarantees that the end user receives the same desktop upon login every time
- Allows corporations to utilize traditional desktop management solutions

This best practice will outline the recommendations for deploying desktops in a VMware Horizon View and Citrix XenDesktop environment.

Windows® Guest Operating System

These best practices apply to all Windows operating systems.

- Disable the screen saver. If a screen saver is required by the company policy, then it may be pushed down through a Group Policy in Active Directory. If this is the case, then choose a blank screen saver.
- Do not configure a sleep timer for these systems.
- Disable any unused ports, such as COM and LPT ports.
- Remove floppy drives from the virtual machine configuration.
- If desired, configure a memory reservation to reduce the amount of storage used by the vSphere vswap file.
- For display properties, choose a basic theme, set the background to a solid color, and disable hardware acceleration.
- Select a high-performance power option.
- Disable the Indexing Service (this may not be true for power users or other users who may search their local system drives often).
- Remove System Restore points
- Turn off system protection on the OS drive.
- Disable any unnecessary services
- Set the sound scheme to No Sound.
- Set visual effects to Adjust for best performance.
- Open Windows Media player and use the default settings.
- Turn off automatic computer maintenance.
- Adjust performance settings for best performance.
- Delete any hidden uninstall folders in C:\Windows.
- Delete all event logs.
- Run Disk Cleanup to remove temporary files, Empty the Recycle Bin, and remove system files and other items that are no longer needed.
- Run Disk Defragmenter to rearrange fragmented data.

Windows 7 Guest Operating Systems

These best practices apply to Windows 7 operating systems.

- Uninstall Tablet PC Components, unless needed.
- Disable IPv6, unless it is needed.
- Use the File System Utility (fsutil) command to disable the setting that keeps track of the last time a file was accessed.
- Change the registry TimeOutValue REG_DWORD in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\Disk to 0x000000be(190).
- With the guest OS shutdown, change the virtual machine video card RAM setting to 128 MB.
- The following Windows 7 Services and Tasks should be disabled.
 - Windows Hibernation
 - Windows Scheduled Disk Defragmentation
 - Windows Registry Backup
 - System Restore
- The following Windows 7 Services and Tasks may be disabled dependent on corporate policy.
 - Windows Update Service - If View Composer recompose is used to install Windows updates and refresh to return OS disks to their original snapshots.
 - Windows Diagnostic Policy Service - If the diagnostic tools are not needed to function on the desktops.
 - Prefetch/Superfetch - if application startup times are acceptable after this feature is disabled.
 - Windows Defender - If other anti-spyware software is installed.
 - Microsoft Feeds Synchronization task - if users do not require automatic RSS feed updates on their desktops.

For more information on optimizing a virtual desktop master image in a VMware Horizon View environment see the [VMware® Horizon View™ Optimization Guide for Windows 7 and Windows 8](#). For more information on optimizing a virtual desktop master image in a Citrix XenDesktop environment see the [Citrix Windows 7 Optimization Guide](#) and [Citrix XenDesktop Planning Guide](#).

VMware Horizon View

Table 5 has information and links for more information on working with Full Clones and Linked Clones in a VMware Horizon View environment with Hitachi Data Systems storage arrays and Hitachi Data Systems compute blades.

Table 4. VMware Horizon View Information

User Type	VM Type	IOPS/Desktop	Link
Heavy Power User	Horizon View Full Clone	30-35	http://www.hds.com/assets/pdf/persistent-vmware-horizon-view-desktops-on-hitachi-ucp-pro.pdf
Knowledge User	Horizon View Linked Clone	25-30	http://www.hds.com/assets/pdf/deploy-hitachi-ucp-select-for-vmware-vsphere-with-vmware-view-heavy-workloads.pdf
Task User	Horizon View Linked Clone	4-7	http://www.hds.com/assets/pdf/deploy-hitachi-ucp-select-vmware-view.pdf

Citrix XenDesktop Linked Clones

Table 6 has information and a link for more information on working with Linked Clones in a Citrix XenDesktop environment with Hitachi Data Systems storage arrays and Hitachi Data Systems compute blades.

Table 5. Linked Clones Information

User Type	VM Type	IOPS/Desktop	Link
Task User	XenDesktop Linked Clone	4-7	http://www.hds.com/assets/pdf/deploy-hitachi-ucp-select-for-citrix-xendesktop-with-vmware-vsphere.pdf

Flex SE VMDKs and Flex SE Linked Clones

SE Sparse Reclaim is a space reclamation technique for shrinking .vmdk files stored on both block and file based storage. Having the capability to reclaim unused space within the VDI guest provides the benefits of linked clone deployments without the need to refresh or recompose. This allows for a more persistent desktop use case by using linked clone pools.

As data is written to and deleted from VMware Horizon View desktops, the VMDKs for the persistent desktops and Linked Clone desktops can grow and take up more space than is needed.

Take this example of SE Sparse Reclaim. Using a VDI Linked Clone desktop, the end user copies a DVD ISO file to the desktop then after the ISO is no longer needed, it is deleted. Without SE Sparse Reclaim, that VDI Linked Clone would still be consuming all the disk space used by the deleted ISO on the enterprise storage system. SE Sparse Reclaim will make that disk space available to the VDI environment again by marking the blocks used by the DVD ISO and shrinking the VMDK file.

For more information on Flex SE VMDKs and Flex SE Linked Clones see [Space Efficient Virtual Disks with VMware Horizon View 5.2 using Hitachi NAS Platform 4080](#).

VDI User Data Redirection

When working with VDI in a VMware Horizon View or Citrix XenDesktop environment it is important to ensure user data and customizations are not stored on the virtual desktop image. If data and customizations are stored on the virtual desktop image, then several challenges can occur.

- The virtual desktop image can grow to an unexpected size causing storage issues.
- In some VDI configurations the end user may not always be assigned the same desktop image at login.
- When the virtual desktop images are recomposed user data and customizations would be lost.
- Virtual desktop images are usually stored on higher tiered and more expensive storage. User data can usually be stored on SATA or near line SAS drives.

In a VMware Horizon View environment it is recommended to use VMware Horizon View Persona Management to ensure that user data is redirected and stored on CIFS infrastructure such as Hitachi Data Ingestor (HDI) backed by Hitachi Content Platform (HCP). For more information on storing VMware Horizon View Persona Management data with HCP and HDI, see [Using Hitachi Data Ingestor and Hitachi Content Platform with VMware Horizon Workspace](#).

In a Citrix XenDesktop environment it is important to use Microsoft Active Directory to implement roaming profiles or use Citrix XenDesktop Provisioning Services to implement Personal vDisks that can be stored on a lower tiered storage.

Resilience and Recoverability Best Practices

vSphere Metro Storage Cluster

A VMware vSphere Metro Storage Cluster architecture on Hitachi storage platforms provides an ideal solution for maximizing availability and uptime by clustering physical datacenters within metro distances. The metro storage cluster solution from Hitachi Data Systems consists of storage systems presenting replicated storage as a single LUN from different geographically distributed sites. This design enables high availability of services by allowing virtual machine migration between sites with no downtime.

A combination of Hitachi software and hardware provides the following key functions to a vSphere infrastructure:

- Host multipathing
- Internal and external storage provisioning
- Synchronous storage replication across metro cluster distances
- Storage failover

These functions work together with VMware vSphere vMotion, vSphere High Availability, and vSphere Dynamic Resource Scheduler to build this reference architecture for a VMware vSphere Metro Storage Cluster. The advanced functionalities found in Virtual Storage Platform fulfill the requirements of a virtual infrastructure and lessen the need for additional hardware that may be required in traditional Metro Storage Cluster solutions.

For information on creating a VMware Metro Storage Cluster see [Deploy VMware vSphere Metro Storage Cluster on Hitachi Virtual Storage Platform](#).

VMware Site Recovery Manager

VMware vCenter Site Recovery Manager is a business continuity and disaster recovery solution that integrates VMware vSphere infrastructures with data replication on storage systems to protect large sites and business critical applications. It does the following:

- Provides centralized management of disaster recovery plans
- Enables non-disruptive testing
- Automates site recovery and migration processes

Using an automated and tested recovery plan, a production data center can do a disaster recovery failover in an organized and proven manner.

Remote data replication is a key function in building out stable and reliable disaster recovery environments. Replicating data to a remote secondary site represents the most effective insurance policy against a catastrophic failure. Although you can perform data replication at the server level, you can perform data replication more effectively within the storage infrastructure.

VMware vCenter Site Recovery Manager integrates tightly with Hitachi storage arrays using the Hitachi Storage Replication Adapter. This use provides centralized management of recovery plans. Tight integration between storage arrays, VMware vCenter, VMware vCenter Site Recovery Manager, and the Hitachi Storage Replication Adapter ensure a coordinated recovery for large, business critical environments.

For more information on implementing VMware Site Recovery Manager with Hitachi storage platforms see [Deploying VMware Site Recovery Manager 5.0 with VMware vSphere 5.0 on Hitachi Virtual Storage Platform](#).

Hitachi Data Protection Suite, Powered by ComVault with Hitachi Thin Image

Using Hitachi Thin Image snapshots with Hitachi storage arrays and controlling these snapshots with Hitachi Data Protection Suite and IntelliSnap provides an efficient and effective solution to address the needs of data protection for VMs.

Hitachi Thin Image hardware snapshots are used to capture images of VMware snapshots, and the VMware snapshots are then quickly removed, eliminating the problems caused by long-term VMware snapshots. Hitachi Data Protection Suite (HDPS) and IntelliSnap automates all the manual steps necessary to create Hitachi Thin Image snapshots, and provides a single management interface for snapshot backup/ restore between Hitachi storage arrays and the VM data.

For more information on using Hitachi Data Protection Suite and IntelliSnap in VMware vSphere Environments see [Using Hitachi Data Protection Suite IntelliSnap to Control Hitachi Thin Image Snapshots in VMware Environments](#).

Hitachi Data Instance Manager

Hitachi Data Instance Manager (HDIM) is a unified, easy-to-use, policy- and workflow-based solution that centralizes multiple data protection capabilities and provides an agentless interface to manage data protection activities for VMware ESXi hosts. It includes backup, continuous data protection, archive, snapshot, deduplication, data security and more, with top-down policy management. HDIM protects data and reduces, controls and manages copy data, or data instances. It replaces or augments legacy data protection point-solutions and eliminates backup windows, improves service levels, reduces backup data volumes, supports compliance requirements and improves application uptime, all while helping to reduce costs.

You can see Hitachi Data Instance Manager's easy to use drag and drop user interface in Figure 5.

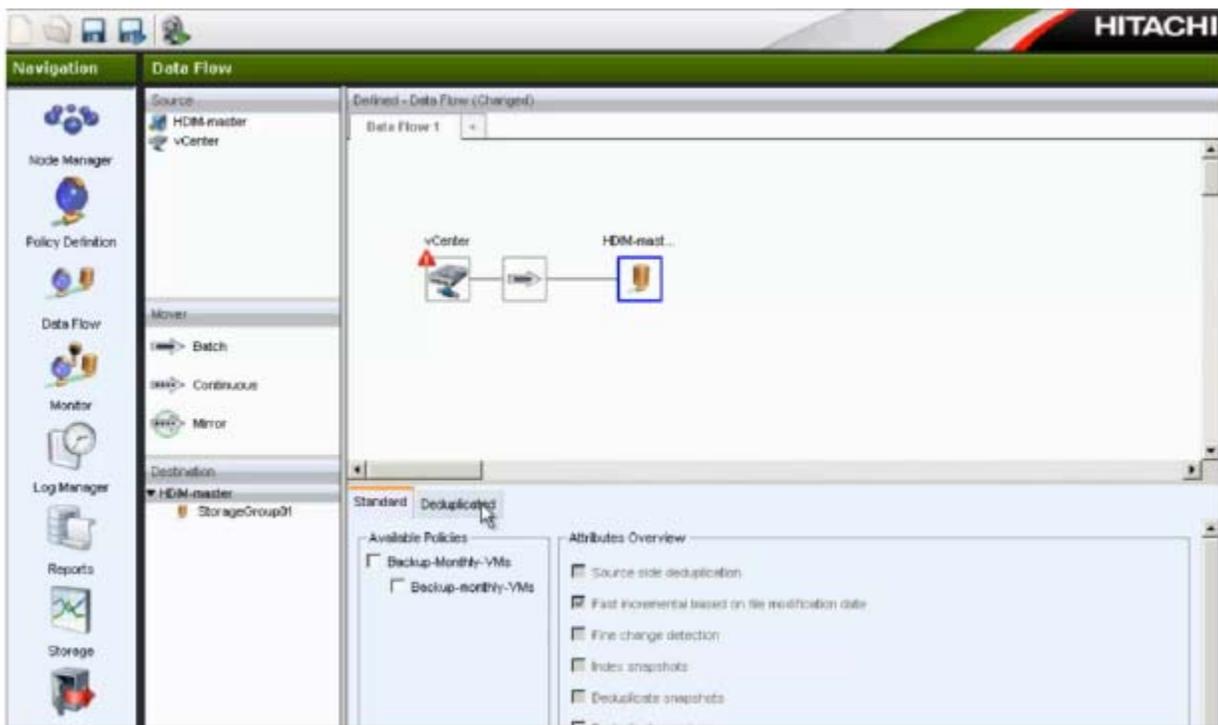


Figure 5

Conclusion

Hitachi Data Systems is an Elite Partner in VMware's Technology Alliance Partner program, a participant in VMware Ready Partner programs for Storage Infrastructure Services, and an embedded OEM partner. Together, Hitachi and VMware are committed to providing innovative, business-enabling, end-to-end virtualization solutions for the data center.

Through Hitachi Data Systems platforms, products and native VMware integrations, Hitachi provides a complete solution for today's rapidly expanding VMware environments.

This Best Practices paper provides guidance to customers either planning to implement or who have already implemented Hitachi Data Systems products with VMware products. This guidance will aid in building a solution that will help to ensure a VMware environment that provides performance, scalability, reliability, usability, resilience and recoverability when paired with Hitachi Data Systems products,

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems [Global Services](#) website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate [Resources](#) website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services [Education](#) website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the [Hitachi Data Systems](#) website.

 **Hitachi Data Systems**



Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

© Hitachi Data Systems Corporation 2016. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. Microsoft, Exchange, SQL Server, and Windows are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

AS-288-01. February 2016.