TECHNICAL PAPER

Protect Hitachi Unified Compute Platform 2000 for VMware vSphere with Hitachi Data Instance Director and Hitachi Virtual Infrastructure Integrator

Lab Validation Report

By Jeff Chen

June 2016

Hitachi Data Systems

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

Protect Hitachi Unified Compute Platform 2000 for VMware vSphere with Hitachi Data Instance Director and Hitachi Virtual Infrastructure Integrator

Lab Validation Report

Hitachi Data Instance Director (HDID) provides business-defined data protection for organizations looking to modernize, simplify and unify their operational recovery, disaster recovery, and long-tern retention operations. HDID is also integrated with Hitachi Virtual Infrastructure Integrator to provide storage-based protection of the VMware vSphere environment.

This lab validation report demonstrates the use of HDID along with Virtual Instance Director to protect Hitachi Unified Compute Platform (UCP) 2000 for VMware vSphere. HDID and Virtual Instance Director can be deployed within UCP 2000 as a single rack data protection solution with options for replicating data to the outside of the rack.

Among many features provided by HDID, this paper focuses on the VMware VADP (VMware vStorage API for Data Protection) backup option and the Virtual Instance Director storage-based snapshot option with the following benefits:

Hitachi Data Instance Director with VMware VADP

- Agentless backup using VMware native API
- Incremental backup that provides a shorter backup window
- Easy to implement and maintain
- Easy to replicate backup data to other destinations

Hitachi Virtual Infrastructure Integrator

- Agentless backup
- No backup window for crash-consistent backup
- Short backup window for application-consistent backup
- Fast recovery
- No network traffic
- Easily scalable

This lab validation report does not provide sizing of storage and virtual machines, but it provides some reference data points and expectations of HDID and Virtual Instance Director backups. Only single rack contained HDID and Virtual Instance Director implementations are referenced in this paper. This paper is intended for storage or data center administrators who are backing up data and recovering data within a VMware vSphere environment.

Note: Testing of this configuration was done in a lab environment. Many factors affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated, test environment that matches your production environment before your production implementation of this solution.

Test Environment Configuration

This section describes the key components and storage configuration used to conduct the test.

Key Hardware Components

Table 1 shows the hardware components.

Table 1. Hardware Components

Hardware	Description	Version	Quantity
UCP 2000	▪ T41S-2U chassis ▪ 2 × fans ▪ 2 × power supplies	Firmware Revision: 3.30.00	1
	▪ T41S-2U compute ▪ 2 Intel Xeon E5-2620 v3 processors, 2.40 GHz ▪ 192 GB memory ▪ Intel 82599 10GigE OCP dual-port card ▪ Emulex LPe12002 dual-port 8 GB HBA card	BIOS Version: S2S_3A14 Emulex HBA Firmware v2.01A10 Intel driver 3.9.58.9101 Emulex HBA driver v10.4.246.0	3
	▪ Hitachi Virtual Storage Platform G400 ▪ Dual controller ▪ 8 × 16 Gb/sec Fibre Channel ports ▪ 64 GB cache memory ▪ 24 × 1.2 TB 10k RPM SAS disks ▪ 12 × 3.2 TB FMD DC2 drives ▪ 2 × NAS modules	83-03-20-40/03	1
	▪ Brocade ICX7450 switch ▪ 24-port 1 GbE management switch	08.0.20c	1
	▪ Brocade VDX 6740 switch ▪ 48-port 10 GbE switch	5.0.1d	2
	▪ Brocade 6505 switch ▪ 24-port 8 GB/sec Fibre Channel	v7.4.1	2

Hitachi Unified Compute Platform for VMware vSphere

offers the following:

- Start small and grow your infrastructure as more workloads are transitioned to a private cloud.
- Simplify management of physical and virtual resources with tight integration into VMware vCenter.
- Improve troubleshooting with physical and virtual infrastructure monitoring and alerting within VMware vCenter.

Unified Compute Platform for VMware vSphere 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

Hitachi Unified Compute Platform 2000 provides the following for your infrastructure:

- Entry-level converged infrastructure solution.
- Flexible, customizable with the ability to grow.
- Pretested, preconfigured, and prebuilt to meet your converged infrastructure needs.

VSP G400 and VSP G600 with NAS Modules

Optional high-performance NAS modules are available in the VSP G400, VSP G600 and VSP G800 systems. These modules help to reduce capital expenditures by providing a single, SAN and NAS storage platform for all workloads in a compact form factor. Operating costs also decline, as VSP G series systems require less power, cooling and space in the data center.

Key Software Components

Table 2 describes the key software components used for testing.

Table 2. Software Components

Software	Version
Hitachi Data Instance Director	5.2.2
Hitachi Virtual Infrastructure Integrator	3.0.45
VMware vSphere	6.0
VMware vCenter Server	6.0
Workload Virtual Machine	Microsoft Windows 2012 R2

Hitachi Data Instance Director

Hitachi Data Instance Director (HDID) software provides modern data protection with business-defined copy data management, enabling the simplified creation and management of complex data protection and retention workflows.

For unified management, HDID delivers a broad set of fully integrated capabilities, both storage-based and host based, that enable a business-driven approach to data protection.

Hitachi Virtual Infrastructure Integrator

Hitachi Virtual Infrastructure Integrator provides a comprehensive data protection and management solution for VMware vSphere environments. It provides you the ability to deliver scalable data protection and management services for your VMware ESXi based virtual infrastructure.

Virtual Infrastructure Integrator simplifies data management with scheduled and instantaneous backup, recovery, and cloning services reducing risks, costs, and administrative overhead. It helps you meet backup and recovery SLAs at VM-level granularity while improving resource utilization.

Virtual Instance Director allows VM administrators to manage application-consistent data protection from the VMware vCenter console, leading to simplified IT operations.

Storage Configuration

Table 3 lists the storage pools and configuration used for this test.

Table 3. Storage Configuration

Pool	RAID Level	Drive Type	Number of Parity Groups	Description
UCP SAN_OS_Boot	RAID10 (2D+2D)	10K SAS	1	Pre-allocated for UCP 2000 ESXi boot LUNs
UCP Management_VMs	RAID10 (2D+2D)	10K SAS	1	Pre-allocated for UCP 2000 management virtual machines including vCenter and data protection management virtual machines
Test_VM_OS	RAID6 (6D+2P)	10K SAS	1	Hosting OS LUNs for test virtual machines
Test_VM_DB	RAID5 (3D+1P)	FMD DC2	1	Hosting data LUNs for test virtual machines
NAS_Compute	RAID5 (3D+1P)	FMD DC2	1	Presenting system drives for NAS modules
HDID_Repository	RAID6 (6D+2P)	10K SAS	2	Allocated for HDID backup repository
V2I_HTI	RAID6 (6D+2P)	10K SAS	1	Allocated for Virtual Instance Director snapshot backup

VSP G400 NAS Module Storage Configuration

Steps to configure the storage for NAS modules are listed below:

1. 16 LUs were created from NAS_Compute pool as system drives for NAS modules
2. From the 16 system drives, one NAS storage pool was created
3. One 4 KB NAS filesystem was created from the storage pool and assigned to an EVS
4. One NFS export was created from the filesystem
5. One NFS datastore was created on ESXi hosts

Virtual Machine Storage Allocation

For HDID, the following two Microsoft Windows 2012 R2 virtual machines were deployed on the UCP 2000 vCenter.

- HDID-Master: The machine that controls the actions of all other nodes on the system.
  - 4 vCPU
  - 8 GB RAM
  - 80 GB OS VMDK
- HDID-Repo: A general-purpose system designated as the recipient of data in a replication configuration.
  - 4 vCPU
  - 16 GB RAM
  - 80 GB OS VMDK, 4 TB repository VMDK

For Virtual Instance Director, one Microsoft Windows 2012 R2 virtual machine was deployed on the UCP 2000 vCenter.

- Virtual Instance Director-Server: This server receives the queries from Virtual Instance Director vCenter plugin and sends talks to the Hitachi storage. To take advantage of Hitachi storage-based snapshots, a command device from the storage needs to be deployed to this machine.
  - 4 vCPU
  - 8 GB RAM
  - 80 GB OS VMDK

Table 4 lists the high level virtual machine storage allocation used.

Table 4. Virtual Machine Storage Allocation

Virtual Machine	Storage Allocation
HDID-Master	▪ Deployed on the default UCP 2000 management datastore
HDID-Repo	▪ OS was deployed on the default UCP 2000 management datastore ▪ VMDK for the backup repository disk was created on a datastore from HDID_Repository pool
Virtual Instance Director-Server	▪ Deployed on the default UCP 2000 management datastore ▪ Command device was also provisioned from the UCP Management_VMs pool
VMFS Test VMs	▪ OS was deployed on a datastore from Test_VM_OS pool ▪ Data VMDKs were deployed evenly on two datastores created from the Test_VM_DB pool
NFS Test VMs	▪ OS and data VMDKs were created from an NFS datastore hosted by VSP G400 NAS modules

Other configurations regarding Virtual Instance Director and HDID deployment are listed below:

- On the Hitachi storage, create a dedicated user for Virtual Instance Director operation.
- If you are using Hitachi storage-based HDID features, create a dedicated user for HDID as well.

Test Methodology

This is the test methodology used. The purpose of the tests is to observe the behavior and differences of a HDID VMware VADP backup and Virtual Instance Director snapshot while a moderate SQL workload is running.

Test Virtual Machine Configuration and Workload

The following virtual machine configuration was used to perform the test:

- 4 vCPU
- 16 GB RAM
- 40 GB OS VMDK Thin
- 200 GB DB VMDK
    - Provisioned as Eager Zeroed Thick for VMFS datastore as followed by best practice
    - Provisioned as Thin for NFS datastore as followed by best practice
    - 60% of disk was filled with random data
    - 20% of disk was actively accessed during the test
- Microsoft Windows 2012 R2 Operating System

For each virtual machine, the SQL Server workload listed in Table 5 was generated against DB VMDK.

Table 5. Workload Definition

Workload	Block Size (KB)	Read Percent	Random Percent	IOPS/Virtual Machine
Microsoft SQL Server	64	66%	100%	27

With this workload, the data change rate became 2 GB hourly or 10% hourly. Eight virtual machines with running workload were used for every backup test.

Test Cases

For HDID VMware VADP backup test, two policies were created to achieve two concurrent backups to increase the backup efficiency. If all eight virtual machines were set in one policy, it would result in only a single stream with eight sequential virtual machine backups. The total backup time can be significantly longer. Table 6 shows the scheduled backup settings.

Table 6. HDID VMware VADP Backup Policy Setting

Number of VADP Policies	2 policies with 4 MVs per policy
Frequency	Every 1 hour
Retention	Keep 1 day

For Virtual Instance Director snapshot backup tests, one backup group was created containing eight virtual machines. With this backup group, the minimal HTI snapshots were created in one snapshot operation. On this test, three LUN snapshots (OS LUN and two DB LUNs) were created for each backup operation. Table 7 shows the scheduled backup settings.

Table 7. Virtual Instance Director Scheduled Group Backup Settings

Number of Groups	1 Group with 8 VMs
Frequency	Every 1 hour
Retention	Keep 1 day

Table 8 lists the test cases and their description.

Table 8. Test Cases

Test Case	Description
Test Case 1 - Measure the backup-window and storage usage for HDID VMware VADP backup on VMFS	In this test case, the eight virtual machines' DB VMDK are deployed evenly on two ESXi hosts with two VMFS datastores. The workload is running for 36 hours during the backup test. The measurement is taken with both quiesce options enabled/disabled.
Test Case 2 - Measure the backup-window and storage usage for HDID VMware VADP backup on NFS	In this test case, the eight virtual machines' DB VMDK are deployed evenly on two ESXi hosts with one NFS datastore. The workload is running for 36 hours during the backup test. The measurement is taken with both quiesce options enabled/disabled.
Test Case 3 - Measure the backup-window and HTI pool capacity usage used by Virtual Instance Director on VMFS	In this test case, the eight virtual machines' DB VMDK are deployed evenly on two ESXi hosts with two VMFS datastores. The workload is running for 36 hours during the backup test. The measurement is taken with both quiesce options enabled/disabled.
Test Case 4 - Create a cloned virtual machine from HDID backup	In this test case, a virtual machine is restored after taking a HDID backup. The timestamp of a restore operation is measured.
Test Case 5 - Restore a virtual machine from a Virtual Instance Director snapshot backup	In this test case, a virtual machine is restored after taking a Hitachi Virtual Infrastructure Integrator snapshot backup. The timestamp of a restore operation is measured.
Test Case 6 - Mount a VMDK snapshot to the original virtual machine using Virtual Instance Director	In this test case, VMDKs are mounted back to the original virtual machine after taking a Hitachi Virtual Infrastructure Integrator snapshot backup. The timestamp of a mount VMDK operation is measured.

Analysis

This section discusses the results and analysis of each test case.

Test Case 1 and 2 - HDID VMware VADP Backup on VMFS and NFS Datastores

Initial backup for VMware VADP takes longer to complete because it requires copying the entire virtual machine data to the repository. Table 9 shows the time that it took to complete the initial full backup and the storage space used in the HDID repository.

Table 9. Initial Full Backup

Datastore Type	Backup Time	Storage Used
VMFS	1 hour 14 minutes	1276 GB
NFS	2 hours 15 minutes	1060 GB

The total storage space used for the virtual machines deployed on VMFS was 1276 GB, and compared to the NFS datastore, it was about 200 GB more consumption. This was because Eager Zeroed Thick format was used for the VMDKs provisioned from VMFS, and the entire VMDK data was backed up. For NFS datastore, Thin provisioning was used, and with Thin format, only consumed space was backed up.

The VMware Change Block Tracking (CBT) feature was utilized for the incremental backup. Table 10 shows the average incremental backup time for VMFS and NFS datastores with the quiesce option enabled or disabled.

Table 10. Incremental Backup

Datastore Type	Quiesce	Backup Time	Storage Used Per Backup
VMFS	ON	6 minutes 31 seconds	31.02 GB
VMFS	OFF	4 minutes 54 seconds	30.9 GB
NFS	ON	8 minutes 34 seconds	30.92 GB
NFS	OFF	7 minutes 31 seconds	30.96 GB

For eight virtual machines with two concurrent backup policies, it took about 5 to 9 minutes to complete hourly incremental backups. With the quiesce option enabled, it took 1 to 2 minutes longer compared to setting the quiesce option to disabled. The average storage space used for each backup was about 31 GB.

Each test was performed for 36 hours with hourly backups, and with a retention of 24 hours. Figure 1 shows the total capacity used for eight virtual machines on VMFS and NFS datastores. Again, backup of virtual machines on VMFS datastore consumed more space because they were provisioned as Eager Zeroed Thick. The capacity usage of the backups from both VMFS and NFS datastore became flat after 24 hours because backup older than 24 hours were deleted.

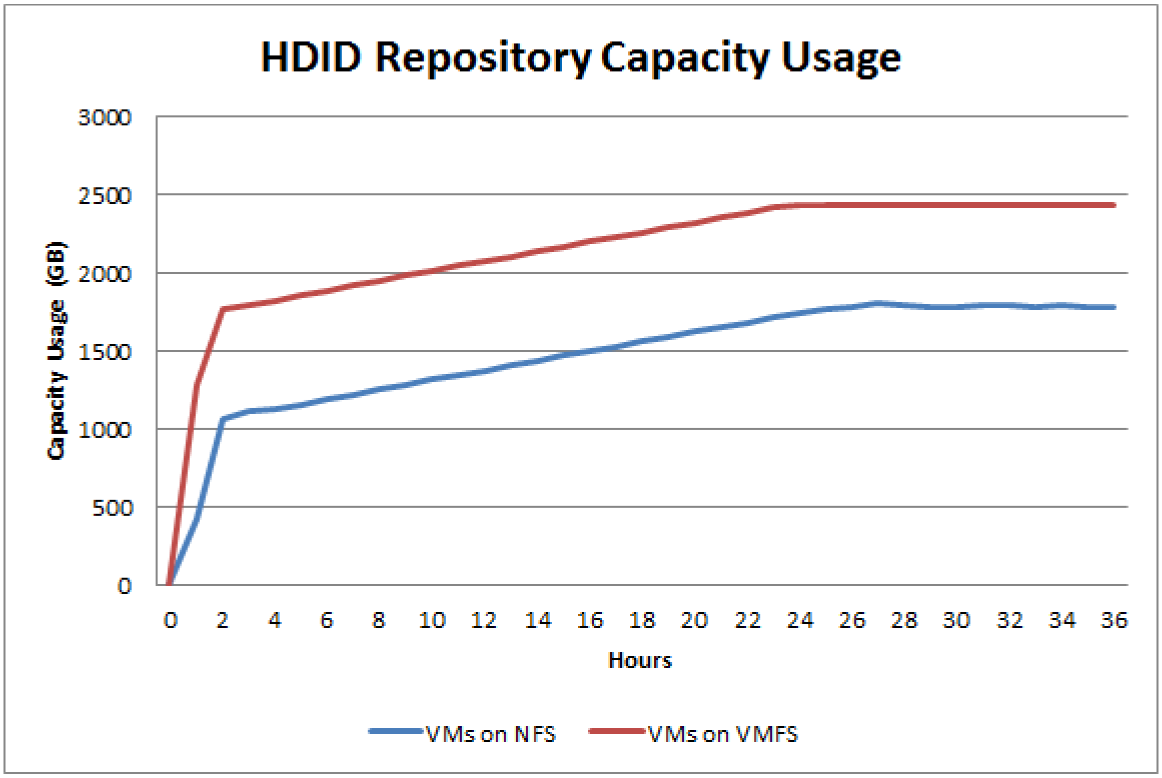

Figure 1

Test Case 3 - Virtual Instance Director Snapshot Backup on VMFS Datastores

The eight virtual machines on VMFS datastores with running workload were backed up hourly by Virtual Instance Director for at least 36 hours. Each backup was kept for 24 hours. Table 11 lists the average backup time with the quiesce option enabled and disabled.

Table 11. Virtual Instance Director Snapshot Backup Time

Quiesce	Backup Time before 24 Hours	Backup Time after 24 Hours
ON	48 seconds	3 minutes 16 seconds
OFF	42 seconds	2 minutes 45 seconds

For this particular test environment, there were not significant differences in terms of backup time with quiesce enabled or disabled. However, it took at least 2 more minutes to delete storage-based HTI snapshots after 24 hours.

Figure 2 shows the total HTI pool capacity usage during this 36 hour test. For this workload setting, the HTI pool capacity increased an average of 56 GB hourly. After 24 hours, it stabilized at 1396 GB as snapshots older than 24 hours were deleted.

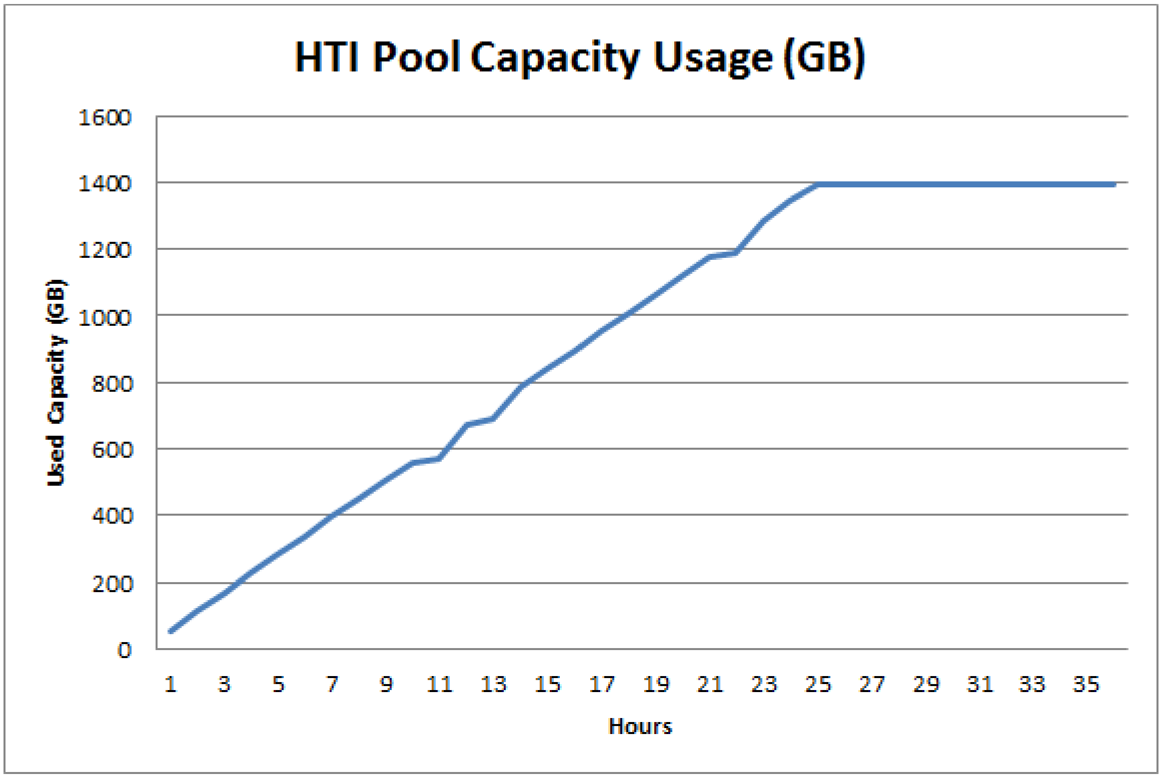

Figure 2

Test Case 4 - Create a Cloned Virtual Machine from HDID Backup

In this test, a virtual machine was restored as a new clone from an HDID backup. The backups were performed from Test Case 1 and Test Case 2 with virtual machines deployed on VMFS and NFS. Table 12 shows the time that it took to restore a virtual machine and the amount of data restored.

Table 12. Restore a Virtual Machine from HDID Backups

Datastore Type	Restore Time	Restored Data
VMFS	37 minutes 5 seconds	213 GB
NFS	21 minutes	133 GB

Test Case 3 - Restore a Virtual Machine Snapshot from a Virtual Instance Director Snapshot Backup

A restore operation can be performed for following reasons:

- A virtual machine is accidentally deleted
- A virtual machine has become corrupted

In this test, a virtual machine was restored from a Hitachi Virtual Infrastructure Integrator snapshot successfully. The average time for a virtual machine to be restored and powered up was 2 minutes and 20 seconds. After powering up the restored virtual machine successfully, storage vMotion was used to migrate it back to the original datastores from the snapshot LUNs.

Test Case 4 - Mount VMDK Snapshots Back to the Original Virtual Machine Using Virtual Instance Director

If files need to be restored, the mount VMDK operation can be performed.

In this test, VMDKs from a snapshot backup were mounted back to the original virtual machine successfully. The average time for this operation was 2 minutes and 55 seconds.

Considerations

With the combination of both HDID and Virtual Instance Director, broader data protection options can be achieved on the VMware virtualized environment. Observation and consideration for each solution is discussed in this section.

HDID VMware VADP

Consider following when implementing the HDID VMware VADP solution.

- Assign HDID server IP subnet to be the same as the target ESXi hosts to minimize network routing.
- During the backup, the maximum throughput of 200 MB/sec was observed for a single virtual machine. Consider the following to help increase total throughput.
  - Use multiple concurrent backup streams at same time
  - On the HDID repository server, add multiple VMDKs from multiple datastores to create HDID Storage Groups
  - Make sure you have adequate backend disk spindles to support high sequential write I/O from the repository server
- High read I/O was observed on the source datastores during the backup.
  - This might impact your application workload.
  - Avoid scheduling backup operations during peak times
- Long initial full backup was observed.
  - Large virtual machines might require multiple recovery point objective (RPO) periods to complete the initial backup.
  - The full backup is only required once, but plan this carefully.
- With the VMware VADP CBT feature, the backup window for the incremental backup was relatively short and optimized.

Hitachi Virtual Infrastructure Integrator

To protect virtual machines on VMFS datastores, the Hitachi Virtual Infrastructure Integrator utilizes Hitachi Thin Image to take LUN-based snapshots. If you have virtual machines with the same backup policy, place them in the same datastore. Use Hitachi Virtual Infrastructure Integrator to take datastore-level snapshots of all of the virtual machines in the same datastore so they can be captured together at the same time.

Since snapshot operations are offloaded to the storage, they only require a minimal backup-window. This gives you an opportunity to take backups more frequently (shorter RPO). Especially with crash-consistent backup, there is no backup-window and no I/O impact from an application perspective. Refer to the following documentation for more information regarding Virtual Instance Director VMFS performance:

- Data Protection for Hitachi Unified Compute Platform using Hitachi Virtual Infrastructure Integrator

As shown in this paper, the Hitachi Thin Image pool requires more capacity if your application workload is highly random with a frequent backup schedule. Take this into consideration when you do HTI pool planning.

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

AS-510-00 June 2016.