HITACHI
Inspire the Next

WHITE PAPER

Protect Hitachi Unified Compute Platform HC with VMware vSphere and Hitachi Data Instance Director

Tech Note

By Chetan Gabhane

June 2017

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.

Revision History

Revision	Changes	Date
AS-595-00	Initial release	May 12, 2017
AS-595-01	Removed limitation of applicability to UCP HC V240F, typographical errors	June 21, 2017

Table of Contents

Protect Hitachi Unified Compute Platform HC with VMware vSphere and Hitachi Data Instance Director

Tech Note

This tech note demonstrates the use of Hitachi Data Instance Director (HDID) to protect Hitachi Unified Compute Platform HC (UCP HC) in a hyper converged infrastructure. The solution focuses on the VMware vStorage API for Data Protection (VMware VADP) backup option for software-defined storage. Data Instance Director protects a VMware vSphere environment as a 4-node chassis data solution with options for replicating data to outside the chassis.

Hitachi Data Instance Director provides business-defined data protection so you can modernize, simplify and unify your operational recovery, disaster recovery, and long-tern retention operations. HDID provides storage-based protection of the VMware vSphere environment.

Data Instance Director with VMware vStorage API for Data Protection provides the following:

- Agentless backup using the VMware native API
- Incremental backup that provides backup window reduction
- Easy to implement and maintain for a virtualization environment
- Easy to replicate backup data to other destinations or outside of chassis

This tech note does not provide information for the sizing of storage and virtual machines. However, it provides some reference data points and expectations of Data Instance Director. This note only contains Data Instance Director implementations with a 4-node Unified Compute Platform HC chassis.

This paper is intended for use by storage or data center administrators who are backing up data and recovering information in 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

These are the key components and storage configuration used to conduct testing.

Figure 1 shows the logical design for testing.

Figure 1

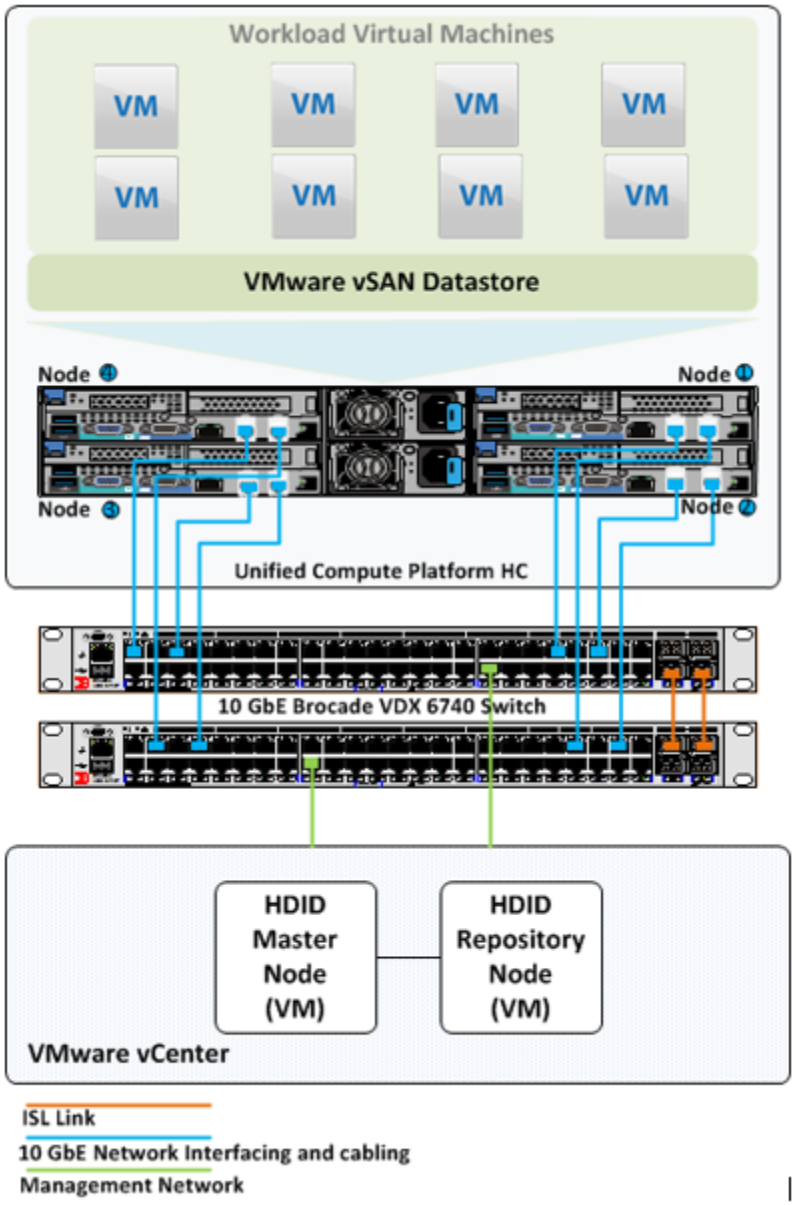

Key Hardware Components

Table 1 describes the key hardware components used for testing.

TABLE 1. HARDWARE COMPONENTS

Hardware	Description	Version	Quantity
Hitachi Unified Compute Platform HC	■ 4-node chassis ■ Intel Xeon processor @2.50 GHz, 12 cores (one processor per host) ■ 2 × 10 GbE NIC ports: ■ NIC options: ■ **RJ-45 connections:** Dual port 10 GbE Base-T Intel X540 OCP mezzanine card ■ **SFP+ connections:** Dual port 10 GbE Intel 82599ES SFP+ OCP mezzanine card ■ 1 × 10/100 Base-T RJ45 port for remote (out-of-band) management ■ 1 × 800 GB SSD for Cache Tier ■ 2 fully redundant power supplies ■ LSI SAS 3008 RAID controller mezzanine card ■ Fault-tolerant virtual SAN datastore ■ SSDs — 4 cache tier ■ HDDs — 9 capacity tier	BIOS Firmware: S2S_3A18 BMC: 3.36.00	1
Brocade VDX6740 switch	■ Top-of-rack switch ■ 10 GbE ports (at least 16 ports) VDX6740	nos 7.0.0.a	2

Hitachi Unified Compute Platform HC

Combining compute, storage, and virtualization into a hyperconverged infrastructure, Hitachi Unified Compute Platform HC (UCP HC) answers challenges of growing demands of faster delivery of business services while facing rising costs managing disparate technology resources. Using VMware Virtual SAN with software from Hitachi Data Systems, this software-defined storage extends the agility and simplicity of the Hitachi Unified Compute Platform family.

Unified Compute Platform HC provides a comprehensive dashboard to view virtual machine. There are health monitors for CPU, memory, storage, and virtual machine usage for the all clusters, including individual appliances and individual nodes.

You need minimal IT experience to deploy, configure, and manage the Unified Compute Platform HC appliance. Leveraging VMware's core products, your administrators can apply existing VMware knowledge, best practices, and processes.

Unified Compute Platform HC 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.

Key Software Components

Table 2 describes the key software components used for testing.

TABLE 2. SOFTWARE COMPONENTS

Software Components	Version
Hitachi Data Instance Director	5.5.2
VMware vSphere	6.5a
VMware vCenter Server	6.5a
Microsoft® Windows Server® (workload virtual machines)	2012 R2
Microsoft SQL Server® (transaction load)	2014
VMware vSAN	6.5

Hitachi Data Instance Director

Hitachi Data Instance Director provides a modern, holistic approach to data protection, recovery, and retention.

Data Instance Director has a unique workflow-based policy engine, presented in a whiteboard-style user interface. It helps you to map copy data management processes to business priorities.

Data Instance Director includes a wide range of fully integrated storage-based and host-based incremental-forever data capture capabilities. These can be combined into complex workflows to automate and simplify copy data management.

VMware vSAN

Seamlessly extending virtualization to storage with an integrated hyper-converged solution that works with your overall VMware environment, VMware vSAN reduces the risk in digital transformation by using existing tools, skillsets, and solutions.

Built by VMware, enjoy the best integration with VMware vSphere features with vSAN. Discover the flexibility to expand with other VMware SDDC and multi-cloud offerings as your needs grow. Protect current storage infrastructure investments with the only hyperconverged infrastructure solution built on policy-based management that extends per-virtual machine policies and automated provisioning to modern SAN and NAS storage systems.

VMware vSAN pools on server-attached storage provides a highly resilient shared datastore suitable for any virtualized workload. This includes business critical applications, virtual desktops, remote IT, disaster recovery, or a developmental operations infrastructure.

Hitachi Unified Compute Platform HC Cluster Configuration

The compute nodes on Hitachi Unified Compute Platform HC were configured as follows:

- **Unified Compute Platform HC Cluster —** The test environment used all four VMware ESXi 6.5a nodes for creating eight workload virtual machines by using VMware vSAN automated node with vSAN datastore size of 12.97 TB (Figure 2).
- **Hitachi Data Instance Director** — The test environment placed management and administration components on a separate infrastructure cluster.

Figure 2

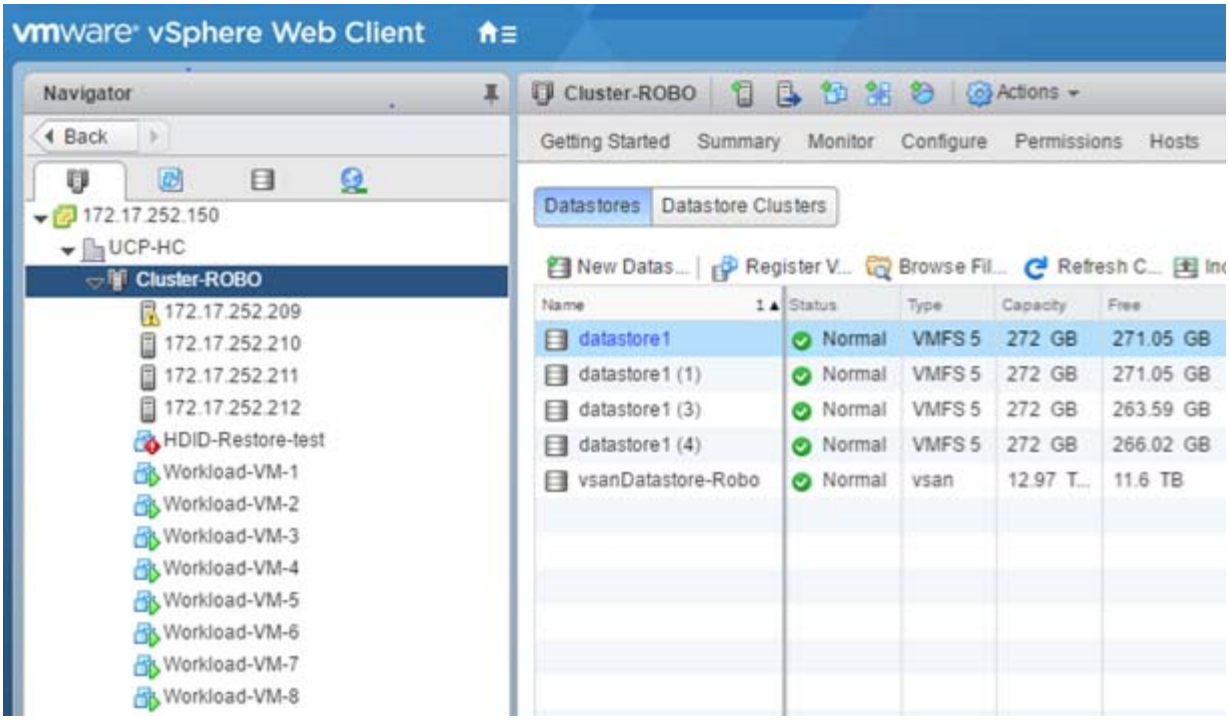

Virtual Machine Storage Allocation

For Hitachi Data Instance Director, the following two virtual machines running Microsoft Windows Server 2012 R2 were deployed on VMware vCenter in Hitachi Unified Compute Platform HC.

- **HDID-Master** — This machine controls the actions of all other nodes on the system.
  - 4 vCPU
  - 8 GB RAM
  - 80 GB OS VMDK
- **HDID-Repo** — This is 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

    The VMDK repository can be on a different VMware vSAN datastore or on external-storage based datastore, such as a Fibre Channel, NFS, or VVol-based.

Table 3 on page 6 lists the high level virtual machine storage allocation used during testing.

TABLE 3. VIRTUAL MACHINE STORAGE ALLOCATION

Virtual Machine	Storage Allocation
HDID-Master	Deployed on the separate vCenter infrastructure management datastore
HDID-Repo	Deployed on the separate vCenter infrastructure management datastore
vSAN Test VMs	OS and VMDKs both were deployed on a vSAN datastore

Best Practices

Consider the following best practices when you implement this VMware VADP solution with Hitachi Data Instance Director:

- Assign the Hitachi Data Instance Director server IP subnet to be the same as the target VMware ESXi hosts to minimize network routing.
- During the backup, the maximum throughput of 252 MB/s was observed for a single virtual machine. Consider the following to help increase total throughput:
  - Use multiple concurrent backup streams.
  - On the Hitachi Data Instance Director repository server, add multiple VMDKs from multiple datastores to create Hitachi Data Instance Director 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 workload times.
- Long initial full backup was observed.
  - Large virtual machines might require multiple recovery point objective (RPO) periods to complete the initial backup.
  - While the full backup is only required once, plan this backup carefully.

Test Methodology

This is the test methodology used to observe the behavior of a Hitachi Data Instance Director environment on Hitachi Unified Compute Platform HC. The environment performed a VMware VADP backup while running a moderate workload of Microsoft SQL Server transactions.

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 a VMware vSAN datastore, following best practice
  - 60% of disk was filled with random data
  - 20% of disk was actively accessed during the test
- Microsoft Windows Server 2012 R2 operating system
- Microsoft SQL Server

For each virtual machine, the Microsoft SQL Server workload listed in Table 4 was generated against DB VMDK.

TABLE 4. WORKLOAD DEFINITION

Workload	Block Size (KB)	Read Percent	Random Percent	IOPS/Virtual Machine
Microsoft SQL Server	64	66%	100%	32

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

Test Cases

For the VMware VADP backup test using Hitachi Data Instance Director, 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 with one policy.

Table 5 shows the scheduled backup settings.

TABLE 5. VMWARE VADP BACKUP POLICY SETTINGS IN HITACHI DATA INSTANCE DIRECTOR

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

Table 6 lists the test cases and their description.

TABLE 6. TEST CASES

Test Case	Description
Test Case 1 — Measure the backup-window and storage usage for the VMware VADP backup using Hitachi Data Instance Director on a VMware vSAN datastore.	Deploy the eight virtual machine's DB VMDK evenly on two VMware ESXi hosts with VMware vSAN datastores. The workload runs for 36 hours during the backup test. Take the measurement with both quiesce options enabled/disabled. This backup is a full backup, with initial backup and a later incremental backup.
Test Case 2 — Create a cloned virtual machine from the Hitachi Data Instance Director backup	Restore a virtual machine after taking a Hitachi Data Instance Director backup. Measure the timestamp of the restore operation.

Analysis

With Hitachi Data Instance Director, you can achieve broader data protection options on the VMware virtualized environment. With VMware VADP CBT, the backup window for the incremental backup was relatively short and optimized.

Test Results

These are the results of each test case.

Test Case 1 — VMware VADP Backup Using Hitachi Data Instance Director on VMware vSAN

The initial backup for VMware VADP takes longer to complete because it requires copying the entire virtual machine data to the repository.

Table 7 shows the time that it took to complete the initial full backup and the storage space used in the Hitachi Data Instance Director repository.

TABLE 7. INITIAL FULL BACKUP

Datastore Type	Backup Time	Storage Used
VMware vSAN	52 minutes	1920 GB

The total storage space used for the virtual machines deployed on the VMware vSAN datastore was 1920 GB. It was about 200 GB more consumption. This was because the **eager zeroed thick** format was used for the VMDKs provisioned from vSAN, and the entire VMDK data was backed up.

VMware change block tracking (CBT) was utilized for the incremental backup. Table 8 shows the average incremental backup time for the VMware vSAN datastore with the quiesce option enabled or disabled.

TABLE 8. INCREMENTAL BACKUP

Datastore Type	Quiesce	Backup Time	Storage Used Per Backup
VMware vSAN	ON	4 minutes 15 seconds	35.02 GB
VMware vSAN	OFF	2 minutes 35 seconds	34.9 GB

For eight virtual machines with two concurrent backup policies, it took about 2 to 5 minutes to complete hourly incremental backups.

With the quiesce option on, it took 1 to 2 minutes longer when compared to setting the quiesce option to off. The average storage space used for each backup was about 35 GB.

Each test was performed for 36 hours with hourly backups, and with a retention of 24 hours.

Figure 3 shows the total capacity used for eight virtual machines on VMware vSAN datastores. The capacity usage of the backups from both vSAN datastores became flat after 24 hours because the backups older than 24 hours were deleted.

Figure 3

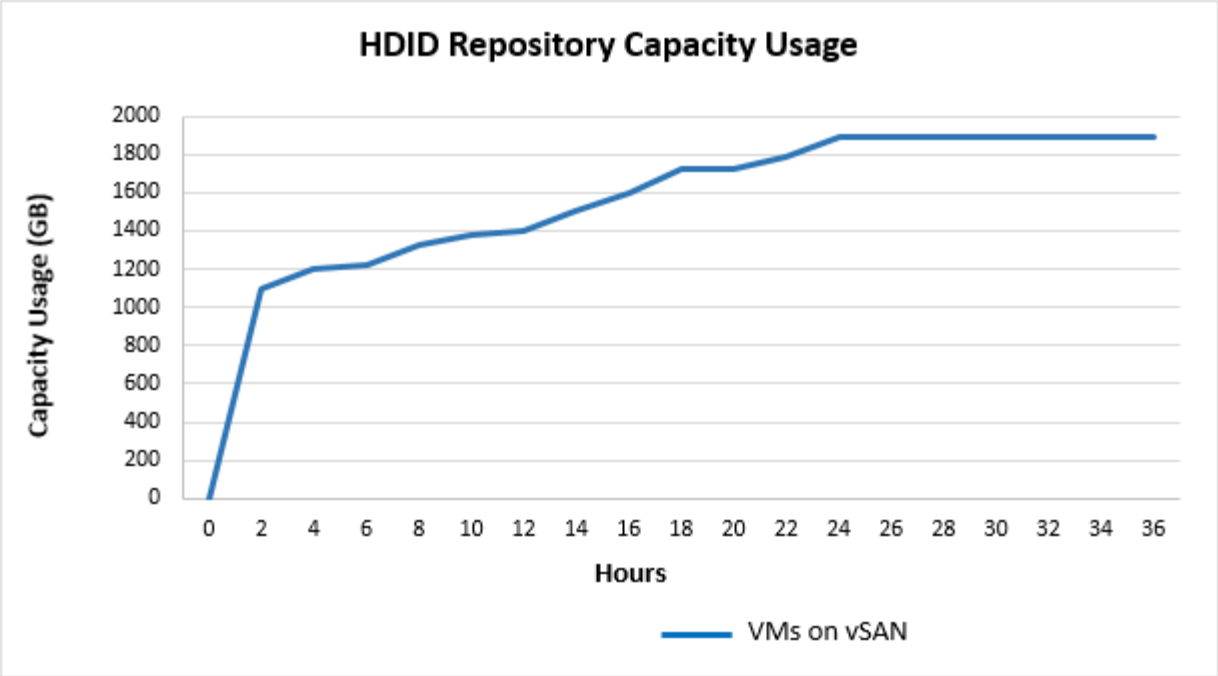

Test Case 2 — Create a Cloned Virtual Machine From Hitachi Data Instance Director Backup

In this test, a virtual machine was restored to the original VMware vSAN datastore as a new clone from a Hitachi Data Instance Director backup. The backups were performed from Test Case 1 with virtual machines deployed on VMware vSAN.

Table 9 shows the time that it took to restore a virtual machine and the amount of data restored.

TABLE 9. RESTORE A VIRTUAL MACHINE FROM HITACHI DATA INSTANCE DIRECTOR BACKUPS

Datastore Type	Restore Time	Restored Data
VMware vSAN	22 minutes 15 seconds	213 GB

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 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 Resources website.

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 Training and Certification 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 95050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 866 374 5822 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-595-01, June 2017.