

Hitachi Data Instance Director on Hitachi Unified Compute Platform HC with VMware vSAN Data Protection Certification

Reference Architecture Guide

By Chetan Gabhane

March 2019

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@HitachiVantara.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
MK-SL-142-00	Initial release	March 8, 2018

Table of Contents

Solution Overview	2
Key Solution Components	4
Key Hardware Components	4
Key Software Components	5
Solution Design	6
Hitachi Unified Compute Platform HC Cluster Configuration	6
Virtual Machine Storage Allocation	7
Best Practices	7
Engineering Validation	8
Test Methodology	8
Analysis	11
Test Results	11

Hitachi Data Instance Director on Hitachi Unified Compute Platform HC with VMware vSAN Data Protection Certification

Reference Architecture Guide

This reference architecture demonstrates the VMware Ready for vSAN certification using 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.

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

These are the benefits of using this Hitachi Vantara-based integrated solution:

- Protect data against disasters, prevent data loss or data integrity issues.
- Enable IT organizations and cloud service providers to store, share, synchronize, protect, preserve, and retrieve file data from a single system. Address future data change and growth.
- Create frequent or temporary backups of the data.

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 reference architecture 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 in one of the following roles who are backing up data and recovering information in a VMware vSphere environment.

- Backup administrator
- Storage administrator
- Data performance analyzer
- Database administrator
- Virtualization administrator
- IT professional with the responsibility of planning and deploying a solution

To use this implementation guide, you need familiarity with the following:

- Hitachi Data Instance Director
- VMware ESXi
- VMware vSAN
- Unified Compute Platform HC Server
- Hyperconverged storage

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.

Solution Overview

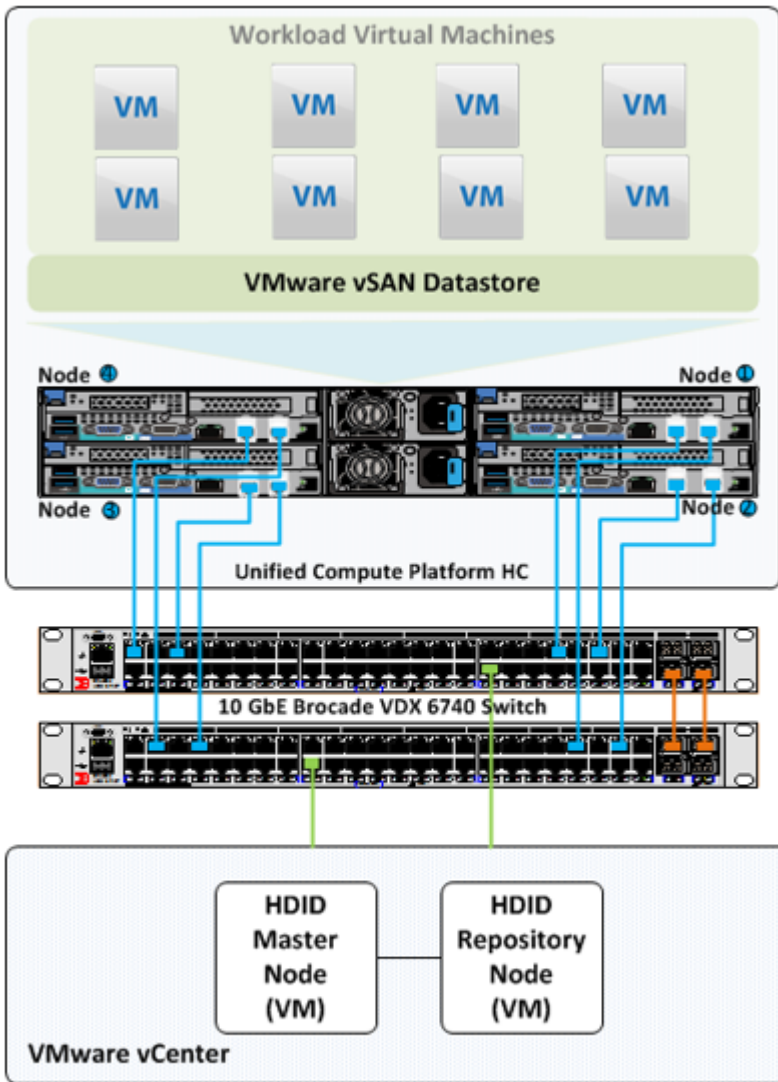
These are the high-level concepts covered by this guide:

- Perform backup of virtual machines (file system data) to the HDID repository store and restore the backed up VM to same vCenter

Use these key components and storage configuration to implement this solution.

Figure 1 on page 3 shows the logical design for testing.

Figure 1



ISL Link

10 GbE Network Interfacing and cabling
Management Network

Key Solution Components

These components are used in this solution.

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	<ul style="list-style-type: none"> ▪ 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 82599EB 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 tiers ▪ SSDs — 12 capacity tiers 	BIOS Version: S2S_3B10.03 BMC Firmware: 3.66.00	1
Brocade VDX6740 switch	<ul style="list-style-type: none"> ▪ 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 Vantara, 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	Version
Hitachi Data Instance Director	6.7 (6.7.1.65371-R6.7)
VMware vSphere	6.5 Build 10390116
VMware vCenter Server	6.5 Build 9451637
Microsoft® Windows Server®	2012 R2
VMware vSAN	6.6

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.

Solution Design

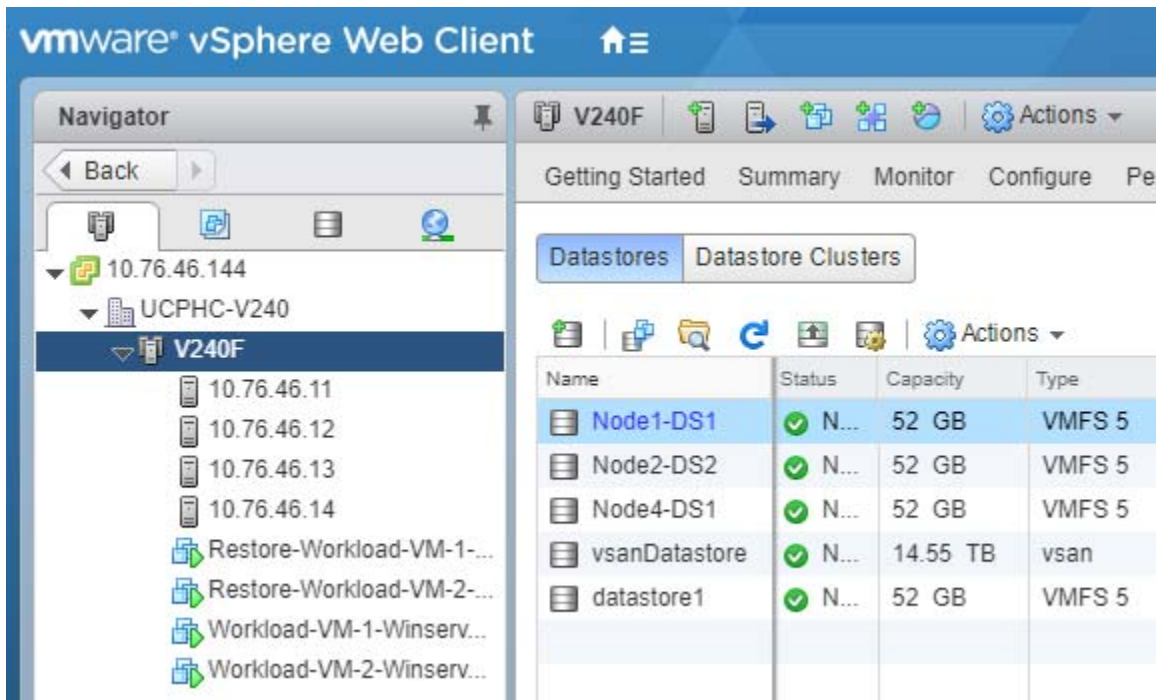
This is the design of this VMware Ready for vSAN certification solution using Hitachi Data Instance Director (HDID) to protect Hitachi Unified Compute Platform HC (UCP HC).

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 6.5 nodes for creating eight workload virtual machines by using VMware vSAN automated node with vSAN datastore size of 14.55 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.
 - 2 vCPU
 - 4 GB RAM
 - 40 GB operating system VMDK
- **HDID-Repo** — This is a general-purpose system designated as the recipient of data in a replication configuration.
 - 2 vCPU
 - 4 GB RAM
 - 40 GB operating system VMDK, 600 GB repository VMDK

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

Table 3 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 VMware vCenter infrastructure management datastore
HDID-Repo	Deployed on the separate vCenter infrastructure management datastore
vSAN Test Virtual Machines	Deployed both the operating system and VMDKs on a vSAN datastore

Best Practices

Follow these 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.
- Because, during the backup, the maximum throughput of 252 MB/s was observed for a single virtual machine during the backup, consider doing the following to help increase your 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.
- Avoid scheduling backup operations during peak workload times. High read I/O was observed on the source datastores during the backup. High read I/O might impact your application workload.
- Large virtual machines might require multiple recovery point objective (RPO) periods to complete the initial backup, as a long initial full backup was observed. While the full backup is only required once, plan this backup carefully.

Engineering Validation

This tells how this design of this VMware Ready for vSAN certification solution using Hitachi Data Instance Director (HDID) to protect Hitachi Unified Compute Platform HC (UCP HC) was validated.

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 thin provisioning for a VMware vSAN datastore, following vSAN default storage policy
 - 60% of disk was filled with random data
 - 40% of disk was free during test
- Microsoft Windows Server 2012 R2 operating system was used

Refer to Table 4 for the workload definition.

TABLE 4. WORKLOAD DEFINITION

Workload	Block Size (KB)	Percent Filled	Percent Unfilled	IOPS/Virtual Machine
Two workload virtual machines, VMware vSAN Datastore, 120 GB filled out of 200 GB using VDbench	64	60%	40%	32

Test cases

For the VMware VADP backup test using Hitachi Data Instance Director, one policy was created to achieve concurrent backups to increase the backup efficiency.

Both virtual machines were set in one policy. It resulted in only a single stream with two sequential virtual machine backups. The total backup time can be significantly longer for first full backup in your environment.

Figure 3 on page 9 and Figure 4 on page 10 shows the for **Policy “vCenter-Backup”** and **Dataflow “vsan-dataflow”** created for the backup in Hitachi Data Instance Director user interface, respectively.

Figure 3

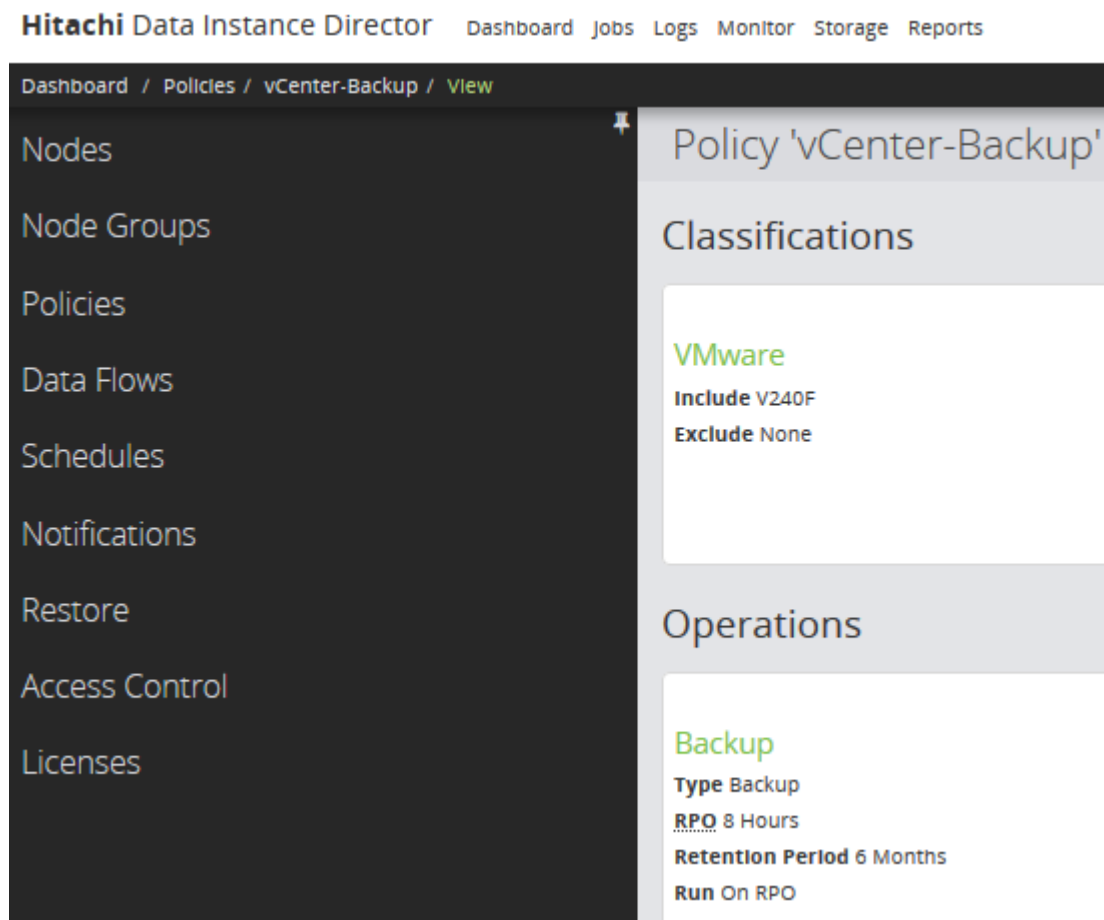


Figure 4

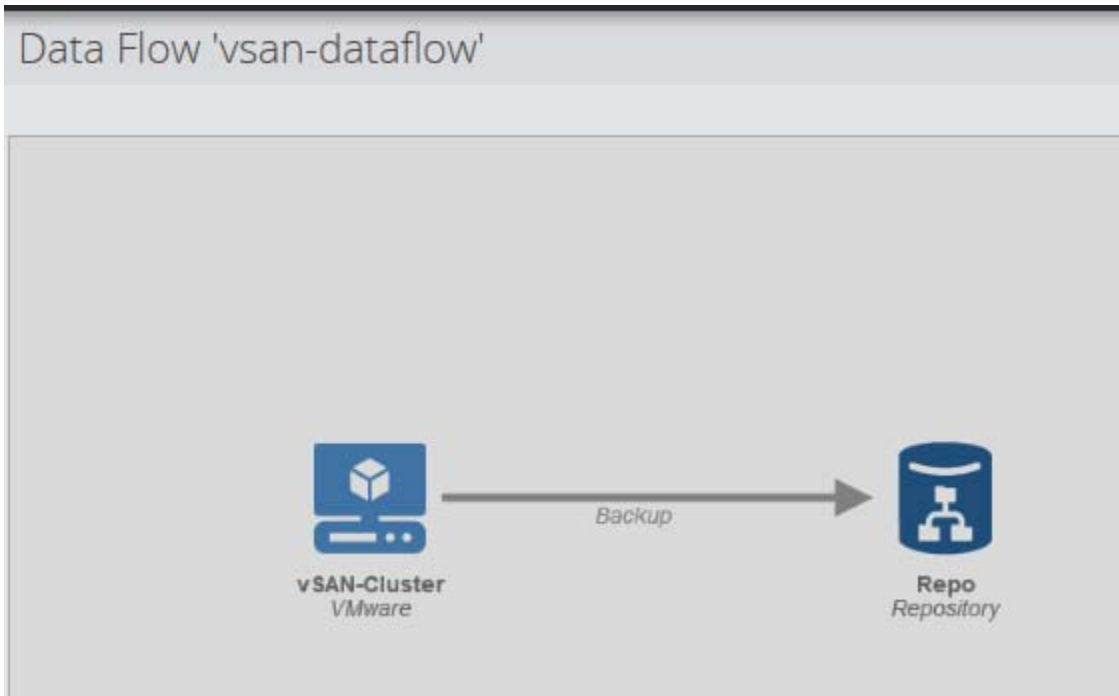


Table 5 shows the scheduled backup settings.

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

VADP Policy	1 policy with 2 virtual machines
Frequency	Once in day
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 two virtual machines with 200 GB VMDK evenly on four node vSAN cluster with VMware vSAN datastores. Backup started with single policy for both workload virtual machines filled with 120 GB out of 200 GB. This backup was a full backup, with initial backup and a later incremental backup.

TABLE 6. TEST CASES (CONTINUED)

Test Case	Description
Test Case 2 — Added 5 GB data to each 200 GB captured incremental backup time	VMDK data was increased during this test by adding 5 GB data by running vdbench on both workload virtual machines and started the backup with same policy.
Test Case 3 — Create a cloned virtual machine from the Hitachi Data Instance Director backup	Restored a virtual machine after taking a Hitachi Data Instance Director backup. Recorded the timestamp of the restore operation.

Analysis

With Hitachi Data Instance Director, you can achieve broader data protection options on a 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

Backup type	Backup Time	Workload Virtual Machines	Storage Used
Initial Full backup	47 min 8 s	2	480 GB

The total storage space used for the virtual machines deployed on the VMware vSAN datastore was 480 GB. During this backup, the entire VMDK data was backed up in the repository virtual machine HDID-Repo.

VMware change block tracking (CBT) was utilized for the incremental backup. For two virtual machines with single backup policy, it took 47 min 8 s to backup 480 GB VMDKs. If there are more workload virtual machines, best practice is to create multiple policies and assign them to workload virtual machines to reduce backup time.

The backup was created once with a retention of 24 hours.

Test Case 2 — Added 5 GB Data to Each 200 GB Captured Incremental Backup Time

Table 8 shows the time that it took to complete the incremental backup and the storage space used in the Hitachi Data Instance Director repository.

TABLE 8. INCREMENTAL BACKUP

Backup type	Backup Time	Workload Virtual Machines	Storage Used
Incremental backup	1 min 15 s	2	480 GB

The total storage space used for the virtual machines deployed on the VMware vSAN datastore was 480 GB. During incremental backup, only the changed blocks of 5 GB from the VMDK were detected and backed up to repository virtual machine. Incremental backup only backed up changed blocks, so backup time is reduced to 1 minute 15 sec. If an incremental backup is not started manually, then it will do an incremental backup on every day at the scheduled time of backup.

Test Case 3 — 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

Backup type	Backup Time	Workload Virtual Machine	Storage Used
Restore	51 min 25 s	2	480 GB

Time taken to restore or clone two virtual machines with single policy is 51 minute 25 seconds.



For More Information

Hitachi Vantara 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.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the [Resources](#) website.

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantara [Training and Certification](#) website.

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

Hitachi Vantara



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

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

© Hitachi Vantara Corporation 2019. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. Microsoft, SQL Server, and Windows Server 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 Vantara.

MK-SL-142-00, March 2019.