
IOmark-VDI-HCI



IOmark

HITACHI
Inspire the Next

Hitachi Advanced Server (HA8000V)

Test Report: VDI-HCI-211030-2

Test Report Date: 30, October 2021



Copyright © 2010-2021 Evaluator Group, Inc. All rights reserved. IOmark-VM, IOmark-VDI, VDI-IOmark, and IOmark are trademarks of Evaluator Group, Inc. in the United States and other jurisdictions.

Table of Contents

Executive Summary	3
Vendor Supplied Product Description	3
IOmark-VDI-HC Test Summary.....	4
IOmark-VDI-HC Results	5
Tested Configuration Details	6
Hyper-Converged System Details	6
Hypervisor Configuration for IOmark-VDI Workload	7
Storage Configuration for IOmark-VDI Workload	8
Configuration Diagram.....	9
Connectivity	9
Certified Configuration Pricing.....	11
Detailed Results	12
Why the Need for IOmark-VDI.....	13
How IOmark-VDI Operates	13
Benchmark Application Workload Set	14
VDI Workload.....	14
VDI Benchmark Parameters	14
VDI Workload Generation.....	14
VDI Workload Details.....	15
Understanding Results.....	15
Benchmark Criteria	15
More Information about IOmark-VDI	15
About Evaluator Group	15

Executive Summary

This document is the official benchmark report for Hitachi Advanced Server (HA8000V), tested as a Hyperconverged system (HCI) using the IOmark-VDI workload.

IOmark is a storage specific workload and benchmark designed to test storage system's performance using a variety of real world, application centric workloads. The IOmark-VDI benchmark is a specific workload, which measures virtual desktop workloads (VDI) run against storage systems. The IOmark-VDI-HCI report certifies that the HCI system is able to support the reported workload, based upon 3rd party compute workload results, combined with IOmark-VDI storage performance results. The IOmark-VDI-HCI results are published after audit and certified approval by IOmark authorized auditors.

The measurement criteria for the IOmark-VDI workload is storage performance. Systems may also be reported as hyperconverged solutions report IOmark-VDI-HCI results and include compute and memory resources in addition to storage, in order to support all requirements for applications tested.

For this test, Hitachi is reporting "IOmark-VDI-HCI" results, which is a Hyperconverged solution including storage along with sufficient CPU and memory required of a hyperconverged solution.

- Hitachi achieved the stated number of IOmark-VDI users for this Hyperconverged solution
- IOmark-VDI-HCI certifies the system supports virtual desktops at the stated price for the HCI infrastructure listed

A full description of the configurations tested along with pricing information is provided in this document. The criteria and performance requirements are as follows:

- For all VDI application workloads:
 - All workloads must reside entirely on the tested system
 - 70% of response times for I/O's must not exceed 30ms
 - The workload replay time must complete within 1 hour and 15 minutes

Vendor Supplied Product Description

Hitachi Advanced Server – HA8000V

Hitachi, Ltd. (hereinafter referred to as Hitachi) has announced new "HA8000V / DL380 Gen10 Plus" and "HA8000V / DL360 Gen10 Plus" of the Hitachi Advanced Server "HA8000V Series" that promotes data utilization with high processing power. These systems support Intel's Xeon® (Series 2R) Scalable processors along with Persistent Memory.

In recent years, various companies have been making efforts to utilize data in order to provide better value to end users. In these efforts, it is required to utilize technologies such as AI and quickly use the large amount of data accumulated every day, so there is an increasing need for improved processing capacity and stable operation in the underlying IT infrastructure. In addition, there are increasing cases of building as a hybrid cloud that combines on-premises and public clouds.

Against this background, the "HA8000V / DL380 Gen10 Plus" are equipped with Intel Corporation CPU and Persistent memory. These systems that supports stable server operation and labor saving in management in an on-premises environment by visualizing device information and regular maintenance by Hitachi maintenance personnel. In addition, "Hitachi pay-as-you-go data infrastructure solution" in line with the service provision concept EverFlex from Hitachi, which realizes "easy-to-introduce price

system”, “flexible use”, and “easy-to-use operation” for customers. This allows customers to use a wide range of rate plans, including monthly billing.

IOmark-VDI-HC Test Summary

For the tested configuration, the following data is provided.

Item	Value
Testing Identifier:	VDI-HCI-211030-2
Product(s):	Hitachi
Test Sponsor:	Hitachi Ltd.
Auditor:	Evaluator Group Inc.

Table 1: Test Identifier Information

Item	Value
IOmark-VDI Version:	Version: IOmark 4.5
Testing Completed:	2021
Equipment Availability:	2020
Audit Certification Date:	30, October 2021
Report Date:	30, October 2021

Table 2: Test Revision and Dates

IOmark-VDI-HC Results

Shown below are the IOmark-VDI-HC results for the system under test. The definition and workload characteristics of the benchmark are provided in Appendix A.

Price information provided below is explained in detail in Table 8. Table 3 below shows an overview of the IOmark-VDI-HC results.

VDI Mode	IOmark-VDI Heavy	Tested Logical Capacity	Total Price	Price / User
Fully Provisioned	384	15.5 TB	\$270,609.00	\$704.71

Table 3: IOmark-VDI-HCI Heavy Price-Performance Results

****Note: Pricing shown is list price and does not include hypervisor or VDI software licenses, pricing details provided in Table 8.***

The “Heavy” workload is measured during a steady-state period during the VDI-HC workload. The measurement period discards the startup and end portions of the workload and measures a 60-minute window during the middle of a 180-minute workload. The response time summary is shown below in Table 4, with further details shown in Figures 3 and 4.

VDI Mode	Test RAID Level	Average Read Resp. Time	Average Write Resp. Time	Average Response Time / VDI User
Fully Provisioned	vSAN Mirroring	0.840 ms	1.911 ms	1.342 ms

Table 4: IOmark-VDI-HCI “Heavy” Passing Result Details

As shown above in Table 4, the Fully Provisioned “Heavy” users had an overall average response time that was lower (better) than the required response time.

Tested Configuration Details

This section covers the connectivity, configuration and pricing information for the system under test.

Hyper-Converged System Details

A single vSAN datastore was created across the cluster which was then utilized by VM's running on the cluster. Access was established to all compute nodes via multi-pathing as noted previously. For this configuration, 1 Compute node was used along with the 3 storage nodes.

Storage System Parameter	Value
Total Servers and Rack Footprint	3 nodes, 6 U – Not including switches
Number of Compute Only Nodes	0
Number of Storage Nodes	3
Hypervisor storage protocol used	vSAN
CPU's	2 per node, Intel 6240R
Memory per node	768 GB - 2133 DDR4
Networking ports	2 @ 10 Gb/s per node
Total capacity available to IOmark-VDI:	Total of 30 TB raw

Table 5: Hyper-Converged Hardware Features

Note, the tested configuration utilized a compute only cluster, along with a hyperconverged storage cluster that provided computing capabilities along with providing the vSAN storage capacity for the entire test configuration.

The CPU and memory necessary to support the reported IOmark-VDI-HCI instances running a Standard VDI workload (via IOmark-VDI standard) required all nodes, each configured with the CPU and memory listed above in Table 5.

Evaluator Group certifies that this configuration is capable of achieving the reported results based upon independent testing, Hitachi test results and other publicly available guidelines for VDI workloads.

Hypervisor Configuration for IOmark-VDI Workload

- Testing certified the use of a vSAN datastore
- A total of 1 datastore was created, with vSAN FS
- Data protection was assigned storage policies per object
 - Setting utilized was mirroring, with checksums (mirrored copies)
- vSAN uses thin provisioning per disk
- **Note:** Deduplication and Compression was not enabled

Detailed hypervisor configuration parameters for the system under test, including connectivity are provided below in Table 6.

Storage System Parameter	Value
vCenter version	VMware 7.0U2 Required for HCI-Mesh
Number of interfaces to the storage system:	2 per node
Connectivity to storage system:	2 @ 10Gb Ethernet ports per node (vSAN interconnect)
Hypervisor storage protocol used:	vSAN
Hypervisor version:	VMware ESXi 7.0U2 required for HCI-Mesh
Thin provisioning:	Yes - enabled
Hypervisor Storage Access:	vSAN over Ethernet
Datastore Filesystem:	vSAN FS version
VASA Interface	Yes (VASA 1.5 / vSAN)
VAAI:	N/A
SATP:	N/A
PSP:	N/A
Total capacity available to IOmark-VDI:	Total of 30 TB raw (15.5 TB with vSAN mirroring)

Table 6: Hypervisor Configuration Parameters

NOTE: Per IOmark requirements, a “write-only” workload is run prior to the actual workload. This pre-writes data to all storage locations referenced during testing. By pre-writing data prior to actual workload testing, there is no write allocation penalty associated with thin provisioning. This also ensures that when reads are performed the storage system reads the media, rather than returning zero’s for unallocated addresses.

Storage Configuration for IOmark-VDI Workload

- A total of 1 vSAN datastore utilizing capacity pooled across all 3 storage nodes
- VMDK's were created for each VDI workload utilizing IOmark tools
- Each VM's VMDK was allocated using "thin provisioning" per Hitachi storage default

Detailed Storage System configuration parameters for the storage system under test, including connectivity is provided below in Table 6.

Storage System Parameter	Value
Storage System firmware	vSAN version 7.0U2
High Availability Access to all LUNs	Yes (vSAN specific)
Total <u>raw</u> capacity of system under test (SUT)	30 TB (across all 3 nodes – not including data compression and deduplication)
Total <u>usable</u> capacity of system under test (SUT)	15.5 TB with vSAN RAID 1 default
Thin provisioning:	Yes
RAID Level(s)	vSAN Mirroring (RAID 1 like mirrored copies)
Raid Pools	2 vSAN disk groups per host
Total Cache Capacity:	600 GB per vSAN disk group, per host (vSAN limit)
Read Cache:	N/A
Write Cache:	600 GB * 1 DG * 3 hosts = 1,800 GB total
VAAI Features Enabled:	vSAN / N/A
- Block Zero	No
- Full Copy	No
- UNMAP	No
- Thin Stun	No
- HW Locking	No
Automated tiering within the storage system:	No
Deduplication or compression of data:	No
Storage system clones / writeable snapshots:	No, not used for testing
Type of storage system clone:	N/A
Storage Media Utilized:	-
- SSD's (Cache)	800 GB SAS SSD per vSAN disk group
- SSD's (Capacity Tier only, Cache noted above)	3.84 TB SATA SSD
- HDD	N/A

Table 7: Storage System Configuration Parameters

Configuration Diagram

The logical data layout of the test configuration is shown below in Figure 1.

A single vSAN datastore was created across the cluster which was then utilized by VM's running on the cluster. Access was established to all compute nodes via multi-pathing as noted previously. For this configuration, 1 Compute node was used along with the 3 storage nodes.

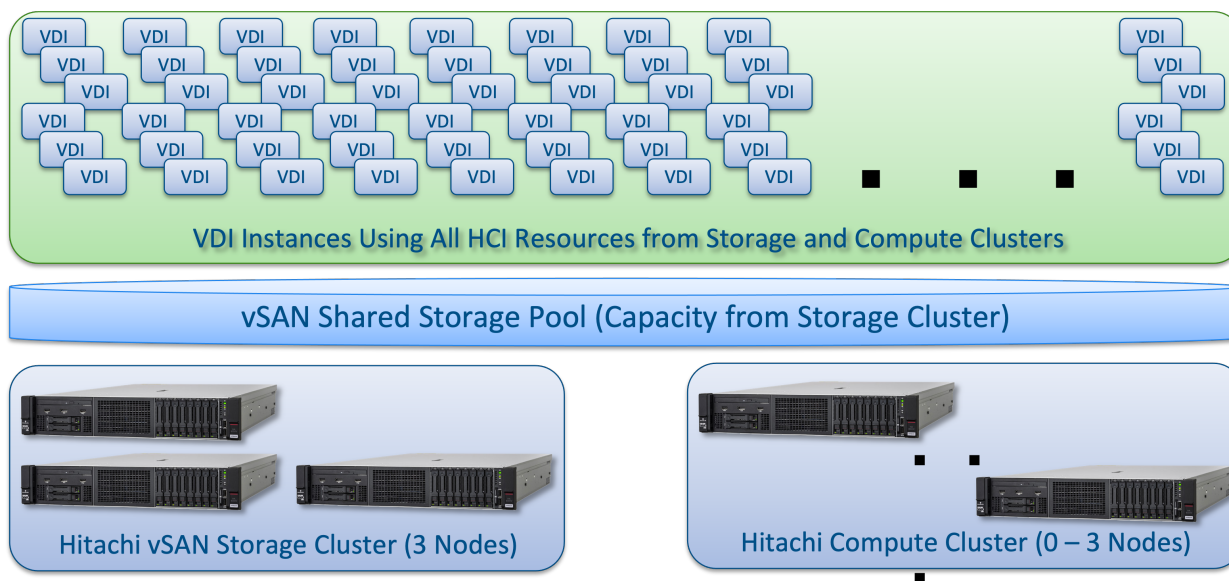


Figure 1: Logical System Configuration

Connectivity

The storage connectivity was 10 GbE for VMware management and data. The test setup utilized 2 x 10GbE network setup, with each node connected to a separate redundant switch for availability per Hitachi best practices.

The configuration connectivity diagram is shown below in Figure 2.

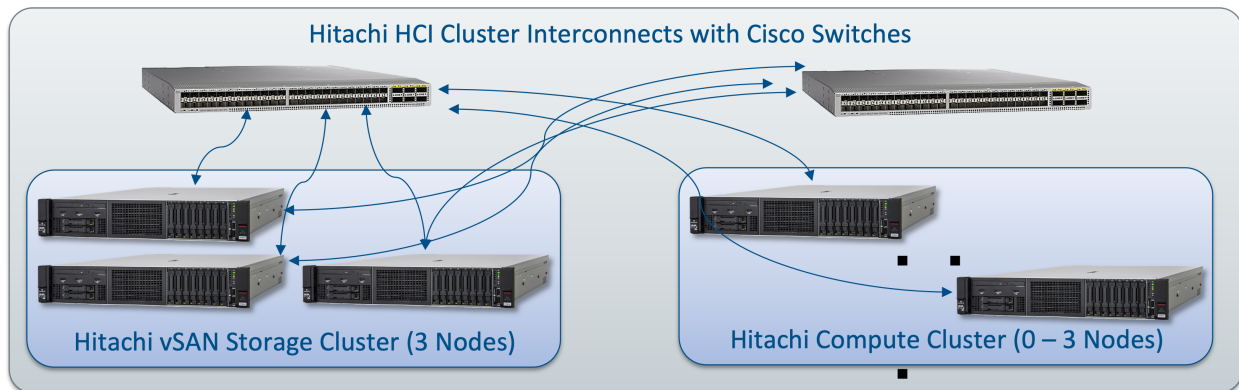


Figure 2: Physical System Connectivity

Certified Configuration Pricing

The certified configuration pricing is shown below, all list prices were provided by Hitachi.

Item	Price / Item	# Items	Extended Price
vSAN no Storage / Node	\$ 63,286.00	3	\$ 189,858.00
vSAN Storage Cost / Node	\$ 26,917.00	3	\$ 80,751.00
Compute Only / Node	\$ 62,004.00	0	\$ -
Total Cost			\$ 270,609.00

Table 8: IOmark-VDI Price Information (list pricing provided by Hitachi)

Detailed Results

IOmark-VDI-HC performance results are measured by application workload. The cumulative distribution function percentages are shown in Figure 3, with response times reported per application in Table 8.




Figure 3: Percentage of Total Response Times at Measured Value

From Figure 3 above, the primary response time(s) of interest are:

- Over 98% of response times were less than 2 ms. for the cluster
- 100% of response times were less than 30 ms. for the cluster

Appendix A - IOmark-VDI Overview

The ability to recreate a known workload is important for comparing a system against potential alternatives. Establishing a reference or benchmark workload enables system vendors as well as resellers and IT users to compare several systems utilizing a known workload.

Specifically, the IOmark-VDI benchmark recreates a storage workload that typically occurs in virtual desktop infrastructure environments. The workload is non-synthetic and recreates several applications that are commonly found in virtualized server environments.

Why the Need for IOmark-VDI

Currently, several application generators have been developed that are able to generate VDI workloads. However, there is no standard reference configuration, with the primary focus is on the server infrastructure. There are no existing benchmark workloads focusing on storage and storage system performance while running VDI applications.

By establishing a set of standard applications and capturing their I/O streams, it is possible to recreate application-based storage workloads for these complex environments. IOmark-VDI is designed utilizing these concepts, and as such is the first benchmark designed to accurately generate application workloads for storage systems, enabling direct comparison of storage system configurations and their ability to support a specific number of applications.

How IOmark-VDI Operates

IOmark-VDI uses the concept of workload replay. I/O streams are captured from actual running applications and then “replayed” so that the exact sequence and I/O commands are issued. This allows the creation of a workload that is indistinguishable from an actual workload to the system under test, while being reproducible and requiring fewer resources. Additionally, the test environment is less expensive, easier and faster to create since actual applications are not required. Because CPU and memory are not consumed running applications, a much higher I/O workload may be generated with a set of server resources than is possible using native applications. This ratio is typically 10:1, but may vary.

In order to scale up the workload on a storage system, additional VDI workloads may be added to the same, or to other physical hosts. The only limitation to the scale of the test is the physical infrastructure supporting the workload. Sufficient, CPU, memory and I/O capabilities must be available to run additional workload sets.

Unlike artificial workload generation tools, IOmark-VDI recreates accurate read vs. write and random vs. sequential I/O requests. Another benefit of IOmark-VDI is the fact that it creates accurate access patterns, thus enabling storage cache algorithms to work properly.

Finally, IOmark-VDI maintains an accurate ratio of performance to capacity as workloads are scaled, ensuring that storage performance is measured with respect to storage capacity accurately. As a result, IOmark-VDI maintains an accurate ratio of I/O to capacity, producing results applicable to IT users.

Benchmark Application Workload Set

VDI Workload

1. View steady state operation
 - a. Heavy Worker Profile – Average / VDI User
 - i. 12.52 iops. / User
 - ii. 1.06 MBps / User
 - b. Standard Worker Profile – Average / VDI User
 - i. 6.26 iops. / User
 - ii. 0.53 MBps / User
2. Benchmark Criteria:
 - 70% of I/O response times must not exceed 30ms
 - All storage utilized must reside on/within the storage system under test

VDI Benchmark Parameters

- Operating System disk size is 20 GB (thinly provisioned)
- All user sessions were running Windows 7 as their guest OS
- No user data disk utilized
- VMware Linked clones may be utilized (as noted)
- Storage linked clones may be utilized (as noted)
- Heavy Profile:
 - The workload is non synthetic, actual I/O patterns are issued based on application capture
 - The size of I/O's is variable, ranging from 512, up to 2 MB transfers based on application
- Standard User Profile:
 - The workload is non synthetic, actual I/O patterns are issued as captured
 - Rates are 50% of "Heavy" user profile
 - The size of I/O's is variable, ranging from 512, up to 2 MB transfers

VDI Workload Generation

The workload generator used to generate the VDI workload was VMware View Planner. This application workload generator controlled running the 8 listed applications above, in a Windows 7 64 bit OS environment, running as a guest VM in a hypervisor environment.

VDI Workload Details

The specific applications comprising a VDI workload set are detailed below in Table 8.

Application	Storage Capacity / Instance
Guest OS (Microsoft Win 10 64bit)	20 GB
MS Office (Word, Excel, PowerPoint and Outlook)	N/A
MS Internet Explorer	N/A
Adobe Acrobat Reader	N/A
Windows Media Server	N/A
Win7 zip	N/A
Total VDI Guest Environment	Total = 30 GB

Table 8: IOmark-VDI Guest Application Overview

The total capacity required for each VDI user is approximately 30 GB of logical capacity. The capacity required for linked clone users is 4 GB, and the capacity required for fully provisioned users is 21 GB without data deduplication or compression.

Understanding Results

IOmark-VDI produces results indicating the response time of a storage system given a particular workload. Based on established criteria, these results indicate how many VDI sessions are supported by a specific storage configuration with a maximum allowed response time. The report is audited for accuracy and issued by Evaluator Group, Inc., an independent storage analyst firm.

Benchmark Criteria

IOmark has established the benchmark criteria for the IOmark-VDI workload. The performance requirements are established as follows:

- For all application workloads:
 - 70% of response times for I/O's must not exceed 30ms
 - All storage must reside on the storage system under test
 - The replay time must complete within 1 hour and 15 seconds for each 1 hour workload

More Information about IOmark-VDI

For more information about the IOmark benchmark, a theory of operations guide, published results and more, visit the official website at <http://www.iomark.org>. Some content is restricted to registered users, so please register on the site to obtain all available information and the latest results.

About Evaluator Group

Evaluator Group Inc. is a technology research and advisory company covering Information Management, Storage and Systems. Executives and IT Managers use us daily to make informed decisions to architect and purchase systems supporting their digital data. We get beyond the technology landscape by defining requirements and knowing the products in-depth along with the intricacies that dictate long-term successful strategies. www.evaluatorgroup.com @evaluator_group