

IOmark-VM



NetApp HCI

Test Report: VM-HC-190802-a

Test Report Date: 2, August 2019



Copyright © 2010-2019 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-VM-HC Test Summary	4
IOmark-VM-HC Results.....	5
Tested Configuration Details	6
Hyper-Converged System Details	6
Hypervisor Configuration for IOmark-VM Workload	7
Storage Configuration for IOmark-VM Workload	8
Configuration Diagram	9
Connectivity.....	9
Tested Configuration Pricing	11
Detailed Results.....	12
IOmark-VM Measurements and Use.....	13
How IOmark-VM Operates.....	14
Benchmark Application Workload Set.....	14
Workload Details	15
Understanding Results	15
Benchmark Criteria.....	16
More Information about IOmark-VM.....	16
About Evaluator Group.....	16

Executive Summary

This document is the official benchmark report for NetApp HCI, tested as a hyper-converged system using the IOmark-VM 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-VM benchmark is a specific workload, which measures Server Virtualization workloads (VMs) run against storage systems. Results are published after audit and certified approval by IOmark authorized auditors.

The measurement criteria for IOmark-VM is storage performance, with the restriction that all storage workloads must be supported by the tested system. For IOmark-VM, CPU and memory considerations, are not tested or considered by the workload.

Systems tested as hyperconverged solutions report IOmark-VM-HC results and include compute and memory resources in addition to storage.

- NetApp HCI achieved the 2nd highest number of IOmark-VM-HC for hyperconverged systems
- IOmark-VM-HC certifies entire HCI system for 1,440 vm applications at \$454.86 / IOmark-VM-HC
 - Configuration: 18 compute nodes + 5 storage nodes
 - System total = \$655,000.00 (\$655,000 / 1,440 = \$454.86)

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 application workloads:
 - All workloads must reside entirely on the tested hyper-converged system
 - Workloads are scaled in sets of 8 workloads
 - 70% of response times for I/Os must not exceed 20ms
 - The average response time for each application type must not exceed 20ms
 - The execution time must complete within 1 hour and 15 seconds for each workload
- For hypervisor operations:
 - Clone, deploy, boot, software upgrade, VM deletion
 - Storage migration (aka Storage vMotion) between storage volumes
- For IOmark-HC benchmarks
 - The system CPU and memory must be sufficient to run the equivalent applications
 - The storage subsystem is measured by IOmark, with system CPU and memory compared to relevant published performance metrics

Vendor Supplied Product Description

NetApp HCI

NetApp HCI is designed for enterprise environments that require the ability to scale storage performance and capacity independently from computing, in order to match application requirements efficiently. NetApp HCI's architecture provides enterprise capabilities along with NetApp Data Fabric components to extend usage, data protection and deployment options.

- **Independent Scale-Out** – Scale storage capacity independently from compute performance

- **Data Efficiency** – In-line deduplication, compression and thin provisioning increase storage efficiency by 5 – 10x
- **Storage Capacity** – From 11.5 TB - 1.8 PB raw capacity (50 TB - 5 PB+ usable with data reduction)
- **Clustering** – Compute scales to 64 nodes per cluster, storage nodes scales from 4 – 40 nodes
- **Quality of Service** – Integrated QoS provides ability to control I/O to isolate applications
- **Data Resiliency** – Dual redundant copies of data distributed to all nodes, automated drive rebuilds
- **Data Protection** – Native snapshot-based backup and restore functionality to object storage via S3 or SWIFT compatible API
- **DR & Replication** – Synchronous, asynchronous and snapshot replication locally and between remote clusters
- **Availability** – Automated failover and failback available between a cluster and up to four other clusters
- **Data Security** – Encryption with 256-bit AES for environments requiring data at rest protection
- **Connectivity** – Storage connectivity via iSCSI, Fibre Channel, VVol's, and container native storage via NetApp Trident

IOmark-VM-HC Test Summary

For the tested configuration, the following data is provided.

Item	Value
Testing Identifier:	VM-HCI-190802-a
Product(s):	5 x NetApp H410S – storage nodes 18 x NetApp H410C compute nodes
Test Sponsor:	NetApp
Auditor:	Evaluator Group Inc.

Table 1: Test Identifier Information

Item	Value
IOmark-VM Version:	Version: IOmark-VM 4.0.1
Testing Completed:	June 2019
Equipment Availability:	March 2019
Audit Certification Date:	31, July 2019
Report Date:	1, August 2019

Table 2: Test Revision and Dates

IOmark-VM-HC Results

Shown below are the IOmark-VM-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-VM-HC results.

IOmark-VM-HC Total VM's	IOmark-VM-HC Response Avg.	Available Capacity (w/ Dedupe)	Used Capacity (w/ Dedupe)	Tested RAID Level	Total List Price	IOmark-VM-HC \$ / VM
1,440	2.35 ms	460 TB (thin + data redux)	6.8 TB	Double-Helix	\$655,000	\$454.86 / VM

Table 3: IOmark-VM-HC Result Details

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

The total number of IOmark-VM virtual machines supported is shown above in Table 3, based on the IOmark-VM workload sets shown in Table 4 below. Each application set consists of 8 virtual machines, thus 180 application sets yields 1,440 VM's reported. The IOmark-VM workload may be used for IOmark-VM-HC configurations.

The VMware vCenter Server™ operation values are also shown below, with two components being reported. The "Clone and Deploy" portion of the workload creates a clone from a specific VM template, starts the VM and then upgrades its version of VMware tools installed. The reported value indicates how many operation cycles were completed during the 1-hour test run. Similarly, the storage vMotion value reported indicates how many migration cycles were completed during the 1-hour test run. A combined score is calculated, known as the "Hypervisor Workload Score," which is the ratio of reported results to the minimum required results. The minimum numbers of vCenter operations for passing the test are 6 clone and deploy and 3 storage vMotion operations respectively for configurations supporting 21 IOmark-VM sets or more.

Details of passing results shown below in Table 4:

IOmark-VM-HC Sets	Read Resp. Average	Write Resp. Average	# vCenter Clone and Deploy	# vCenter storage vMotion	Hypervisor Workload Score (1 - inf.)
180	3.77 ms	1.31 ms	7	7	2.1

Table 4: IOmark-VM-HC Passing Result Details

Tested Configuration Details

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

Hyper-Converged System Details

Detailed hardware features for the system under test are provided below in Table 5.

Hardware Features	Value
Rack Footprint*	18 NetApp Compute Nodes + 5 NetApp HCI Data Nodes nodes = 14 U (10U + 4U)
Number of Nodes per Appliance	H410 chassis supports 4 nodes per 2U rack unit
Number of Drives per Appliance	Cache: 8 GB NVRAM card / storage node = 40GB total Capacity: 6 x 2 TB SSD / data node = 30 total
CPUs (for VM workload)	2 per node → 36 sockets total for 18 compute nodes
Memory (for VM workload)	768 GB per node → 13.8 TB for 18 compute nodes
Networking Ports (10/25 GbE)	(2 per Compute Node + 2 per storage node)

Table 5: Hyper-Converged Hardware Features

*** Note:** The CPU and memory reported above are for the compute nodes only and does not include the CPU and memory utilized by the storage nodes.

The tested configuration utilized multiple servers with VMware vSphere ESXi software for the hyperconverged compute environment, utilizing storage provided by the 5 NetApp H410S storage nodes running Element OS.

The CPU and memory necessary to support the reported 1,440 IOmark-VM's, running the VMmark 2.5 workload required 18 compute nodes, using 2 x Intel Xeon Scalable 6138 processors w/ 768 GB of DRAM. Evaluator Group certifies that this configuration is capable of achieving the reported results, based upon testing and other publicly reported VMmark 2.5 results available.

With these guidelines, the tested Hyper-Converged system achieved the storage performance required and has sufficient computing resources to achieve the stated results.

Hypervisor Configuration for IOmark-VM Workload

- Testing certified both traditional iSCSI LUNs and VMware VVols
- A single VVol datastore was created for the IOmark-VM's certified + 3 iSCSI LUNs
- RAID level was assigned using NetApp HCI default of RAID 1 (Double Helix)
- QoS settings were applied on a per VM basis based upon policies
 - All workloads were assigned "default" IO rate using QoS
 - Note: Use of QoS was not required to achieve passing IOmark-VM results
- NetApp HCI utilizes thin provisioned disks with data deduplication and compression

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

Storage System Parameter	Value
vCenter version	VMware vSphere™ ESXi 6.5 (8307201)
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:	Both VMware VVols and iSCSI
Hypervisor version:	VMware ESXi 6.5 (5969303)
Thin provisioning:	Yes - enabled
Hypervisor Storage Access:	Ethernet
Datastore Filesystem:	Both VMware VVols and VMFS 6
VAAI:	Yes
SATP:	VMW_SATP_DEFAULT_AA
PSP:	VMW_PSP_RR
Total capacity of system allocated to IOmark-VM:	5 Nodes: 45 TB raw

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-VM Workload

- A single VVol datastore was created using the pooled capacity across all nodes
- Three iSCSI datastores were utilized for administrative VMs and as clone and vMotion targets
- Policy-based management was used to allocate VMDK's to each VM for each workload
- Each VM's VMDK was allocated using "thin provisioning" per vSAN 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	Element OS 11.1
High Availability Access to all LUNs	Yes (active / active)
Total <u>raw</u> capacity of system under test (SUT)	45 TB (across all 5 nodes – no deduplication)
Total <u>usable</u> capacity of system under test (SUT)	460 TB (including thin, dedupe and compression)
Thin provisioning:	Yes
RAID Level(s)	Double Helix (RAID 1 like mirrored copies)
Total Cache Capacity:	8 GB x 5 storage nodes = 40 GB total
Read Cache:	N/A
Write Cache:	8 GB NVDIMM per storage node
VAAI Features Enabled:	Yes
- Block Zero	Yes
- Full Copy	Yes
- UNMAP	Yes
- Thin Stun	Yes
- HW Locking	Yes
Automated tiering within the storage system:	N/A (All Flash)
Deduplication or compression of data:	Both
Storage system clones / writeable snapshots:	Clones available, not used for testing
Type of storage system clone:	N/A
Storage Media Utilized:	-
- SSD's (Capacity Tier only, Cache noted above)	NetApp HCI - 2.0 TB TLC 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.

Both LUNs with VMFS datastores and VMware VVols were utilized for testing, with identical performance results. For simplicity, VVols with VMDK based policy-management for QoS is depicted below. For the vMotion workload, two VMFS datastores residing on iSCSI LUNs were utilized.

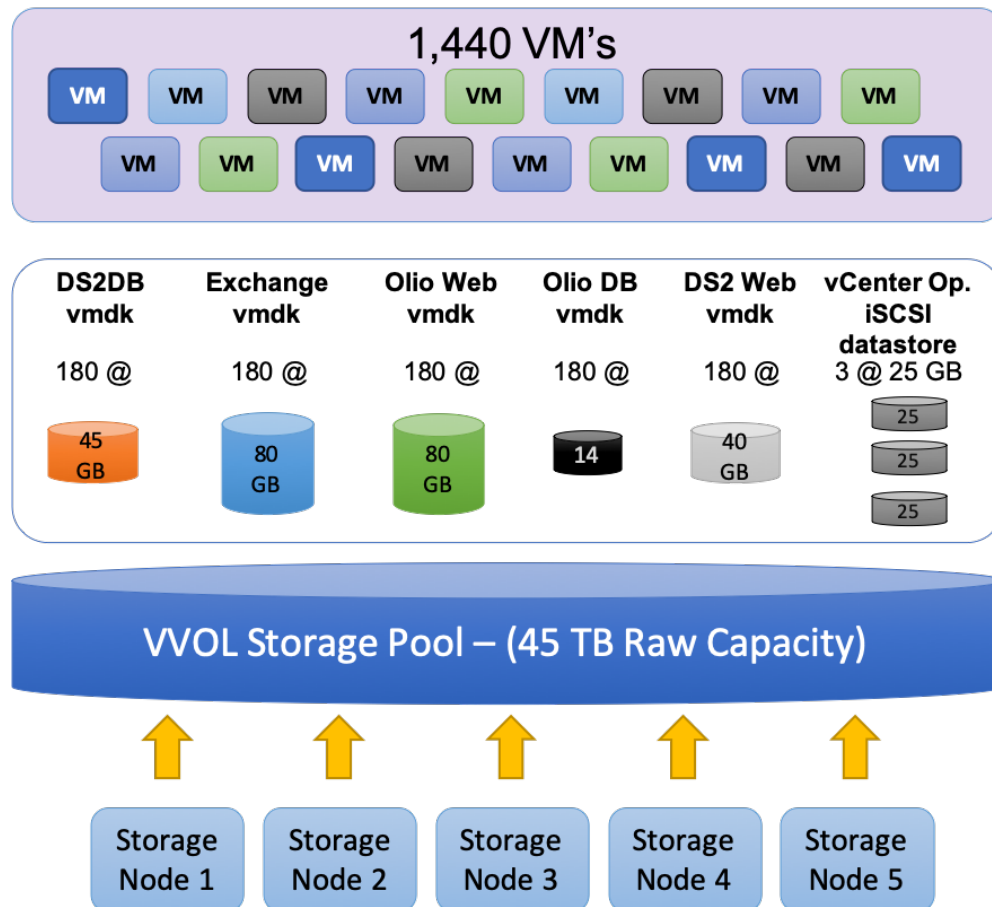


Figure 1: Logical System Configuration

Connectivity

The storage connectivity was 10GbE for both VMware management and the iSCSI LUNs and VVols. NetApp HCI best practices allow for 6 connections per node, or 2 connections per node. The test setup utilized 2 x 10GbE network connection setup, with each node connected to a separate redundant switch for availability per NetApp best practices.

The tested configuration connectivity diagram is shown below in Figure 2.

**18 Server Nodes + 5 Storage Nodes =
46 Total 10 Gb Connections Required**

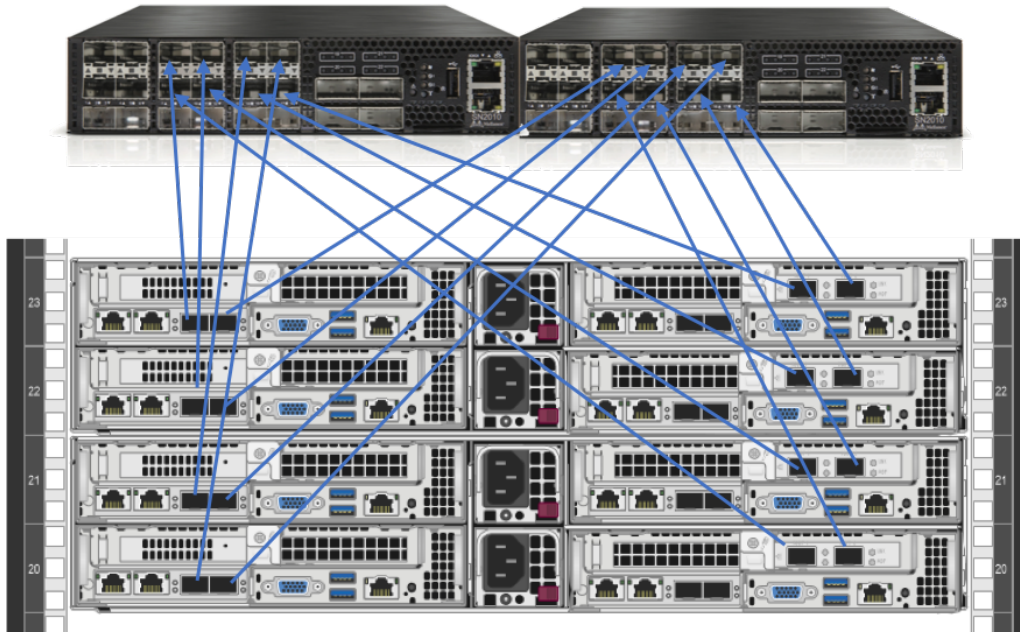


Figure 2: Physical System Connectivity

Certified Configuration Pricing

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

Item	Description	Qty.	List Price
1	NetApp H410C-25025* (2 @ 24 core CPUs, 768 GB DRAM + 2 @ 10 Gb Ethernet)	18 nodes compute	\$15,000
2	NetApp H410S-31110* (5 @ 18 core CPUs, 768 GB DRAM + 2 @ 10 Gb Ethernet)	5 nodes storage	\$77,000
Total	List Price		\$655,000

Table 8: IOmark-VM-HC Price Information (list pricing provided by NetApp)

Detailed Results

IOmark-VM 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.

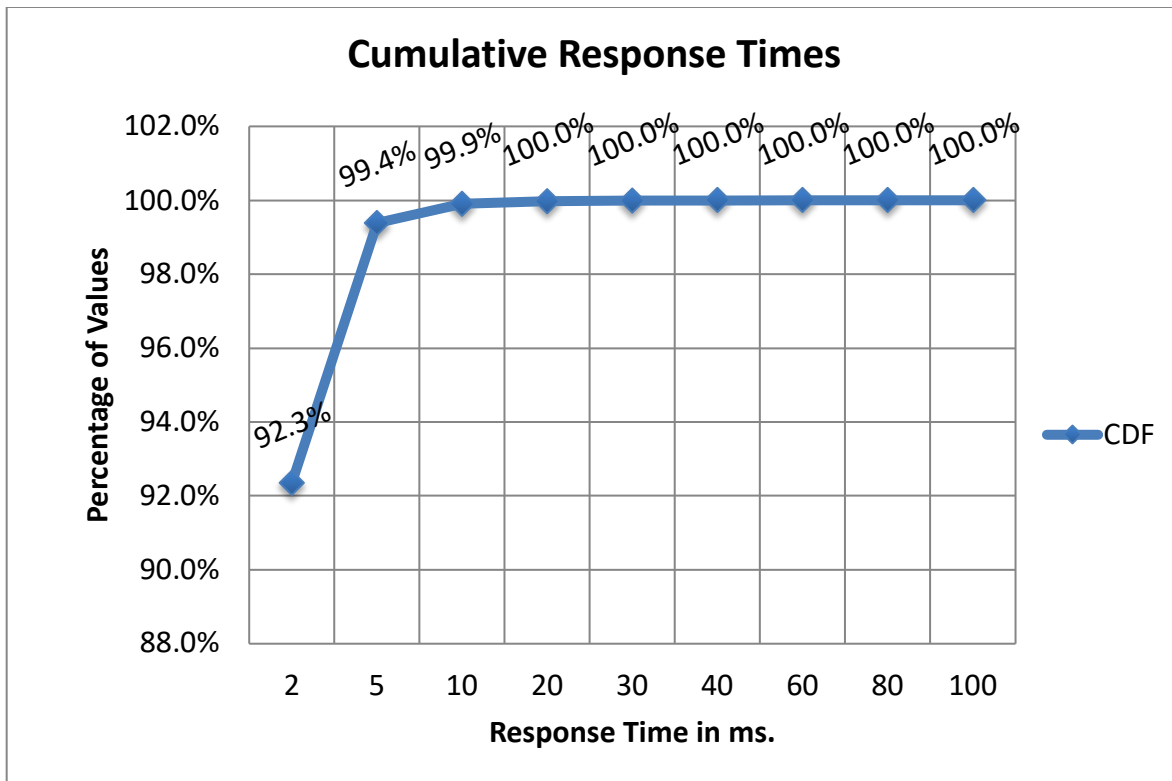


Figure 3: Percentage of Total Response Times at Measured Value

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

- Over 99% of response times were less than 5 ms. for the cluster
- 100% of response times were less than 20 ms. for the cluster

Application Workload	Avg. Response Time	
	Read	Write
DVD Store DB	7.12 ms	1.49 ms
Exchange Mail Server	5.31 ms	1.11 ms
Olio Web Server	2.90 ms	1.31 ms
Olio Database	3.44 ms	0.88 ms
DVD Store Web App 1	0.09 ms	1.77 ms
DVD Store Web App 2	0.09 ms	1.77 ms
DVD Store Web App 3	0.09 ms	1.77 ms
Windows Standby	0.09 ms	1.77 ms

Table 8: Application Workload Response Times

Appendix A - IOmark-VM 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-VM benchmark recreates a storage workload that typically occurs in a virtual infrastructure environment. The workload is non-synthetic and recreates several applications that are commonly found in virtualized server environments.

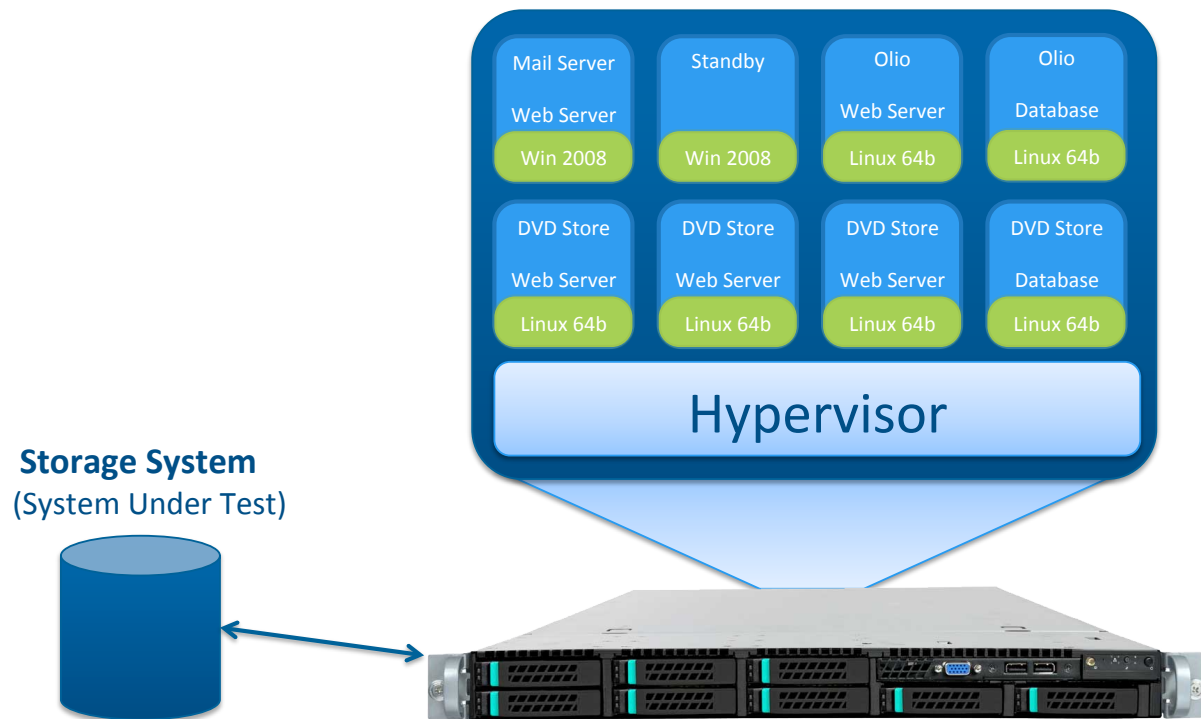


Figure 4: IOmark-VM Conceptual Overview

IOmark-VM Measurements and Use

Datacenters running applications in a virtual infrastructure contain multiple workloads running on a virtualization platform. Often multiple physical servers share the resources of a single storage system providing primary storage for both virtual machine OS and applications.

Currently, several benchmarks have been developed that focus on the server aspects of infrastructure, including the CPU, memory and I/O bandwidth capabilities of the infrastructure. However, there has been no corresponding development of standardized workloads designed to drive storage workloads for these application environments.

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-VM 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.

Additionally, IOmark-VM realizes that a significant impact on storage may occur from administrative functions common in virtual infrastructures. For this reason, several hypervisor-based functions are a part of the IOmark-VM workload. These additional operations include; cloning a virtual machine, booting a VM and updating software, while also migrating a virtual machine from one storage volume to another.

How IOmark-VM Operates

IOmark-VM 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 Figure 4 on the previous page, a single set of applications is depicted running on a single physical host in a virtual infrastructure. In order to scale up the workload on a storage system, additional applications sets 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-VM recreates accurate read vs. write and random vs. sequential I/O requests. Another measurement of IOmark-VM is that it creates accurate access patterns, thus enabling storage cache algorithms to work properly.

Finally, IOmark-VM 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-VM maintains an accurate ratio of I/O to capacity, producing results applicable to IT users.

Benchmark Application Workload Set

A concept utilized for testing multiple applications is that of “Application sets”, also known as “tiles.” A set of 8 applications is run together, along with several common hypervisor infrastructure operations. In order to scale the workload up and place a higher load on the storage system, additional application sets are run. Application sets are always run together for official benchmark results, along with a defined set of infrastructure operations.

The specific applications comprising a workload set are detailed below in Table 10.

Application	Guest OS	Storage Capacity / Instance
Microsoft Exchange 2007	Microsoft Windows Server 2008, Enterprise, 64 bit	80 GB
Olio Database	SuSE Linux Enterprise Server 11, 64bit	14 GB
Olio Web server	SuSE Linux Enterprise 11, 64bit	80 GB
Idle Windows Server	Microsoft Windows Server 2003 SP2 Enterprise Edition, 32-bit	10 GB
DVD Store Database	SuSE Linux Enterprise 11, 64bit	45 GB
DVD Store Web Server 1	SuSE Linux Enterprise 11, 64bit	10 GB
DVD Store Web Server 2	SuSE Linux Enterprise 11, 64bit	10 GB
DVD Store Web Server 3	SuSE Linux Enterprise 11, 64bit	10 GB
Hypervisor Clone & Deploy	N/A - VMware vCenter required	15 GB
Hypervisor Storage Migration	N/A - VMware vCenter required	30 GB
--	--	Total = 305 GB

Table 10: IOmark-VM Application Overview

The total capacity required for each set of applications is approximately 305 GB of capacity. Each additional workload set requires an additional 305 GB of capacity.

Workload Details

The Olio application consists of both a database server, and a web client running on different virtual machines with a pre-loaded data set. For more details on Olio see: <http://incubator.apache.org/olio/>

The DVD application consists of a single database server along with three web clients, each running on a different virtual machine using predefined workload and data set. For more details on the publicly available DVD database application see: <http://linux.dell.com/dvdstore/>

The Exchange server is a Microsoft messaging and email server. Only the server portion of Exchange is recreated in this workload set, with the client workloads not being a part of the I/O, only indirectly through their requests to the messaging server.

The two hypervisor workloads are based on common operations performed in virtual infrastructure environments and require the availability of a VMware vCenter server to perform the operations.

Understanding Results

IOmark-VM produces results indicating the response time of a storage system given a particular workload. Based on established criteria, these results in turn dictate how many total virtual machine

sets are supported by a specific storage configuration and the average response time. The report is audited for accuracy and issued by Evaluator Group, Inc., an independent storage analyst firm.

Note: IOmark-VM response times cannot be directly compared to VMmark response times. IOmark measures response times of individual I/O requests, whereas VMmark measures transaction response times, consisting of multiple I/O operations along with data calculations.

Benchmark Criteria

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

- For all application workloads:
 - Workloads are scaled in sets of 8 workloads
 - 70% of response times for I/O's must not exceed 20ms
 - The average response time for each application 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
- For hypervisor operations:
 - Clone, deploy, boot, software upgrade, VM deletion
 - Storage migration (aka Storage vMotion) between storage volumes

More Information about IOmark-VM

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