


# IOmark-VM

---



vmware®



**VMware vSAN**

**Intel Servers + VMware vSAN Storage SW**

Test Report: VM-HC-170711-a

Test Report Date: 11, July 2017



Copyright © 2010-2017 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-HC Workload.....	6
Storage Configuration for IOmark-VM-HC Workload.....	8
Configuration Diagram .....	9
Connectivity.....	10
Tested Configuration Pricing .....	11
Detailed Results.....	11
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 the tested configuration using Intel® S2600WT server platforms together with VMware vSAN storage software as a hyper-converged system.

IOMark is a storage specific workload and benchmark designed to test storage systems 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.

IOMark-VM-HC is a benchmark that certifies Hyper-Converged systems for virtual server results. The measurement criteria are storage performance, with the restriction that all storage workloads must be supported by the tested Hyper-Converged system. Although there are CPU and memory considerations, these aspects are not tested by the IOMark-VM workload.

The results achieved by Intel hardware with vSAN SW running IOMark-VM is:

- A four-node cluster supported 800 IOMark-VM's at a cost of \$302.04 per VM

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

### Intel Servers with VMware vSAN

Data-center economics are driving new solutions for delivering enterprise class storage. Using Intel based server platforms together with VMware vSAN software-defined storage provides a flexible model for enterprise computing and hyper-converged solutions. With the scale-out architecture, adding additional nodes scales CPU, memory and storage capacity and performance capabilities for the entire system. Server-based storage systems can provide high levels of performance along with economic benefits as shown by the leading price / performance levels achieved during testing.

Intel® Xeon® Scalable processor based systems feature enhancements vs. previous platforms, including:

- A 34% increase in the total Ghz available per chip (64.8 Ghz vs. 48.4 Ghz)
- A 50% increase in memory bandwidth (6 channels vs. 4)

- Enhanced context switching performance, further enhances multi-threaded and virtual apps
- New instructions to improve encryption and other processor enhancements
- New Intel® Volume Management Device (Intel® VMD) enables hot swap of NVMe SSDs from the PCIe bus, while standardized LED management helps provide identification of SSD status

VMware vSAN provides enterprise storage features:

- Hypervisor-embedded storage for ease of management and deployment using existing tools
- Flexible, scale-out and scale-up storage supports adding storage and/or compute on demand
- Enterprise storage with high-availability, synchronous replication with Stretched Clustering
- Enhanced Stretched Clusters with Local Failure Protection
- New vSAN 6.6 enhancements include: improved deduplication and RAID-5 & RAID-6 performance

## IOmark-VM-HC Test Summary

For the tested configuration, the following data is provided.

Item	Value
Testing Identifier:	VM-HC-170711-a
Product(s):	4 x Intel R2208WF with Intel Xeon Platinum 8168 Processors + Intel Media (Intel® DC P4800x Optane) + Intel Flash Media (Intel® DC P4500 3D NAND) + VMware vSAN 6.6
Test Sponsor:	Intel
Auditor:	Evaluator Group Inc.

**Table 1: Test Identifier Information**

Item	Value
IOmark-VM Version:	Version: IOmark-VM 3.8
Testing Completed:	June 2017
Equipment Availability:	July 2017
Audit Certification Date:	11, July 2017
Report Date:	11, July 2017

**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 in this report.

Table 3 below shows an overview of the IOmark-VM-HC results.

IOmark-VM Total VM's	IOmark-VM Response Avg.	Available Capacity (w/ Dedupe)	Used Capacity (w/ Dedupe)	Tested RAID Level	Total List Price	IOmark-VM-HC \$ / VM
800	3.67 ms	58 TB	9.23 TB	RAID 5 & 1	\$241,663.57	\$302.05

**Table 3: IOmark-VM-HC Result Details**

The total number of IOmark-VM-HC 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 100 application sets yields 800 VM's reported.

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 running 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 Sets	Read Resp. Average	Write Resp. Average	# vCenter Clone and Deploy	# vCenter storage vMotion	Hypervisor Workload Score (1 - inf.)
100	1.53 ms	3.90 ms	6	4	2.0

**Table 4: IOmark-VM 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 server hardware features for the system under test are provided below in Table 5.

Hardware Features	Value
Rack Footprint	8U (2U for per R2208WF system)
Number of Nodes per Appliance	1 compute/storage node (4 node cluster)
Number of Drives per Appliance	Cache: 2 Intel DC P4800x Optane / system = 8 total Capacity: 4 Intel DC P4500 / system = 16 total
CPUs per system / total	2 CPUs (2.7GHz Intel 8168) per system @ 24 Hyperthreaded Cores each, for a total of 8 CPUs with 192 Hyperthreaded Cores for the cluster
Memory	768 GB / node => (3 TB for 4 node)*
Networking Ports (1/10 GbE)	8 @ 10GbE Ports (2 Per Node)

**Table 5: Hyper-Converged Hardware Features**

- **\* Note:** It was determined that the amount of memory required to run the reported number of IOmark-VM applications was 768 GB / node, or 3 TB for the cluster. This determination was made based on the footnoted references below, which references systems running VMmark 2.5 workloads noted.

The 4-node cluster achieved an IOmark-VM-HC workload of 100 application sets, or 800 IOmark-VMs. Each IOmark-VM application set is equivalent to one VMmark 2.5 tile.

The Intel Xeon 8168 CPUs provides up to 50% greater performance than previous generation E5-2600v4 CPU's, based upon the greater number of cores, and improved architecture<sup>1</sup>. These improvements enable a 4-node system to run 800 applications (or 100 VMmark tiles), with sufficient CPU capacity necessary to support the CPU utilization of the vSAN software.

By comparison, the tested configuration had faster clock and bus speed for the CPU along with the same amount of DRAM. With these guidelines, the tested Hyper-Converged system achieved the storage performance required and has sufficient computing resources to achieve the stated results.

**NOTE: In order to achieve the VMmark workload levels reported, the configuration requires 3 TB of memory for a 4-node cluster in order to allow direct comparisons to VMmark.**

### Hypervisor Configuration for IOmark-VM-HC Workload

- No traditional LUNs or volumes were created, instead a single vSAN datastore was used
- A virtual disk was created for each of the reported IOmark-VM's certified (320 or 480)
- Data layout and RAID level was determined on a per VMDK basis based upon policies applied

<sup>1</sup> Intel documentation: [Intel-Xeon-Scalable-Processors-ARK](#)

- RAID-1 mirroring policy was utilized for DS2DB and MAILS application disks
- RAID-5 policy was used for remaining application disks
- Dedupe and Compression were enabled for the vSAN datastore
- vSAN default policy of using Checksums was utilized (e.g. checksums enabled)
- VMware vSAN default storage policy utilizes thin provisioned disks

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

Storage System Parameter	Value
Hypervisor	VMware vSphere™ ESXi 6.5.0d
Number of interfaces to the storage system:	2 Per Node (10 total)
Connectivity to storage system:	1 @ 10Gb Ethernet / node (vSAN interconnect)
Hypervisor storage protocol used:	vSAN (Proprietary IP Protocol)
Hypervisor version:	VMware ESXi 6.50d
Thin provisioning:	Utilized vSAN default storage policy
Hypervisor Storage Access:	vSAN datastore
Datastore Filesystem:	vSAN version 5.0
VAAI:	N/A
SATP:	vSAN proprietary
PSP:	vSAN proprietary
Total capacity of system allocated to IOmark-VM:	4 Node: 29 TB raw (vSAN dedupe enabled)

**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 zeros for unallocated addresses.

## Storage Configuration for IOmark-VM-HC Workload

- A single vSAN datastore was created using the pooled capacity across all nodes
- Policy-based management was used to allocate VMDK's to each VM for each workload
- A policy of "RAID-5" was chosen for Olio DB, Olio Web, the standby and 3 DVD Web apps
- A policy of "RAID-1" was chosen for the Exchange Mail and DVD store database apps
- Each VM's VMDK was allocated using "thin provisioning" per vSAN default policy

Detailed Storage System configuration parameters for the storage system under test, including connectivity is provided below in Table 7.

Storage System Parameter	Value
Storage System firmware	vSAN 6.6
High Availability Access to all LUNs	Yes (active / active)
Total <u>raw</u> capacity of system under test (SUT)	29 TB (across all 4 nodes + dedupe enabled)
Total <u>usable</u> capacity of system under test (SUT)	Depends upon RAID and data reduction levels
Disk groups	2 vSAN disk group per node
Thin provisioning:	Yes
RAID Level(s)	RAID 5, 1 (used for different workloads)
Total Cache Capacity:	3 TB (vSAN configuration dependent)
Read Cache:	N/A (Not used in all flash configurations)
Write Cache:	750 GB / node Optane device (3.0 TB total)
VAAI Features Enabled:	Yes- Standard with vSAN
- Block Zero	Yes
- Full Copy	Yes
- HW Locking	Yes
Automated tiering within the storage system:	N/A
Data Deduplication:	Enabled (Note: requires compression)
Data Compression:	Enabled (Note: requires deduplication)
Storage system clones / writeable snapshots:	No
Type of storage system clone:	N/A
Storage Media Utilized:	-
- SSD's (Note: Capacity Tier only, Cache noted above)	Each disk group used 2 x 2 TB Intel P4500 = 29 TB physical usable for 4 node cluster (capacity tier)
- HDD	NA

**Table 7: Storage System Configuration Parameters**

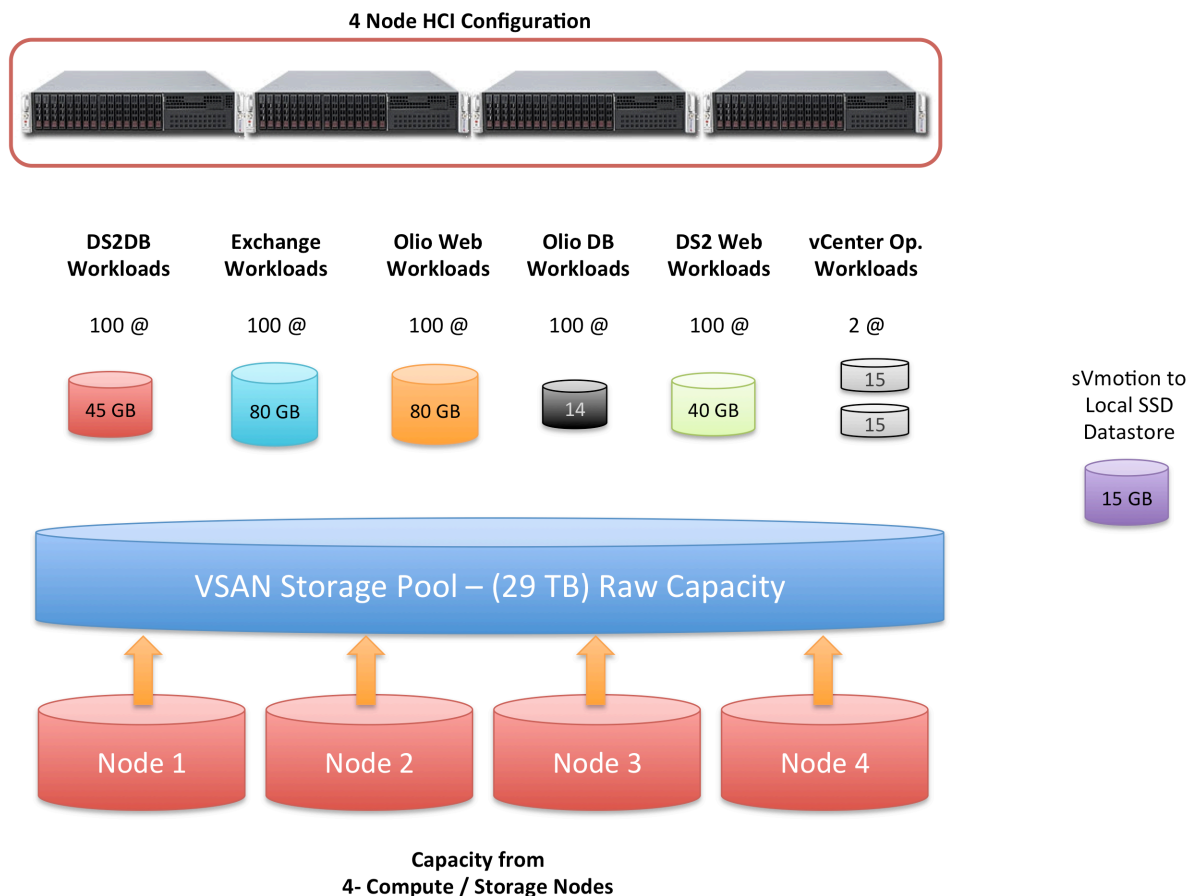


## Configuration Diagram

The logical data layout of the test configuration is shown below in Figure 1. Since vSAN utilizes policy-based storage management, individual LUNs and datastores were not utilized. Instead, individual virtual disks were assigned directly to VM's as required for each workload. The VMDK's for each VM's workload all were allocated from the same vSAN Datastore capacity pool created by vSAN across the 4 storage and compute nodes.

**Note:** Given the available servers and drive configurations, vSAN best practices suggest 2 or more disk groups per node to deliver optimal performance and availability.

**Note:** Since vSAN leverages a single datastore and policies to adjust per-VM capabilities, storage vMotion within that vSAN datastore is not applicable. vSAN supports storage vMotion to external storage systems. In order to provide a valid comparison to other reported IOmark-VM results a local SCSI SSD device was utilized as a target for storage vMotion operations. In order to create the same level of vMotion I/O, two storage vMotion operations occurred in parallel, one moving the VM from the local datastore to vSAN, and another moving the VM from vSAN to the local datastore. The number of storage vMotion operations reported in Table 4 is ½ the number that were actually performed. That is, the bidirectional storage vMotion operations were counted only as a single operation. In this way the results are directly comparable to other tested IOmark-VM results.



**Figure 1: Logical System Configuration**

## Connectivity

The storage connectivity was a single 40GbE port for VMware vSAN. Each node used 1 x 40GbE links to an Ethernet switch, for a total of 4 connections for the 4-node configuration. Testing did not utilize a redundant HA configuration, although production deployment assumes HA network infrastructure.

**Note:** The priced configuration does not include the cost of the 40 GbE card, due to the fact that the measured bandwidth never exceeded 65% of a 10 GbE link at any time (650 MB/s). Moreover, Evaluator Group determination is that the on-board 10 GbE connections would meet bandwidth requirements for the tested configuration.

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

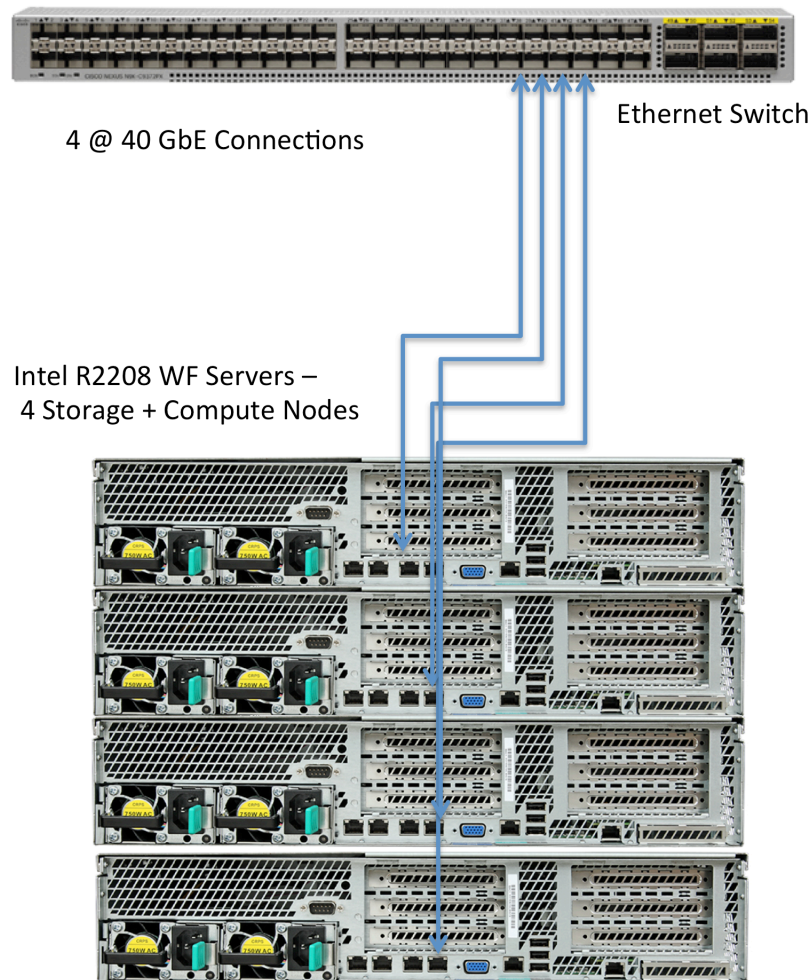


Figure 2: Physical System Connectivity

## Tested Configuration Pricing


Item	Description	Qty.	Ext. List Price
1	Intel R2208WF Standard Servers (2 @ 24 core CPUs, 768 GB DRAM + 2 @ 10 Gb Ethernet)	4 servers	\$76,615.36
2	VMware vSphere Enterprise Plus	8 sockets	\$27,960.00
3	VMware vSphere vSAN Adv.	8 sockets	\$31,960.00
4	Intel DC P4800x 375 GB Optane Storage Devices	8 media	\$12,160.00
5	Intel DC P4500 2 TB 3D NAND Storage Devices	16 media	\$21,728.00
6	3 Years Support HW (10% Server + 5% storage media / year)	4 servers	\$28,097.81
7	3 Years VMware SW Support	8 sockets	\$43,142.40
<b>Total</b>	<b>List Price HW + SW + 3 year service &amp; support</b>		<b>\$241,663.57</b>

**Table 8: IOmark-VM-HC Price Information (4 Node Configuration)**

**Note:** Support included for all hardware and software including Intel flash media.

## Detailed Results

IOMark-VM performance results are measured by application workload. The eight applications that comprise a workload set are shown below in Table 10, with average response times reported per application type.



**Figure 3: Percentage of Total Response Times at Measured Value**

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

- Over 93% of response times were less than 5 ms. for the cluster
- More than 99% of response times were less than 20 ms. for the cluster


Application Workload	Avg. Response Time	
	Read	Write
DVD Store DB	1.8 ms	2.2 ms
Exchange Mail Server	2.0 ms	2.3 ms
Olio Web Server	1.8 ms	5.3 ms
Olio Database	2.0 ms	4.2 ms
DVD Store Web App 1	0.1 ms	5.5 ms
DVD Store Web App 2	0.1 ms	5.5 ms
DVD Store Web App 3	0.1 ms	5.5 ms
Windows Standby	0.1 ms	5.5 ms

**Table 10: 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 3: 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 3 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](http://www.evaluatorgroup.com) @evaluator\_group*