IOmark-VM

Datrium DVX
Test Report: VM-HC-171024-b
Test Report Date: 27, October 2017

Table of Contents

Executive Summary ... 3

Vendor Supplied Product Description ... 3

IOmark-VM-HC Test Summary ... 4

IOmark-VM-HC Results ... 6

Tested Configuration Details ... 7

- Hyper-Converged System Details ... 7
- Hypervisor Configuration for IOmark-VM-HC Workload ... 8
- Storage Configuration for IOmark-VM-HC Workload ... 9
- Configuration Diagram ... 10
- Connectivity ... 11
- Priced Configuration at U.S. List Price ... 12
- Detailed Results ... 13
- IOmark-VM Measurements and Use ... 14
- How IOmark-VM Operates ... 15
- Benchmark Application Workload Set ... 15
  - Workload Details ... 16
- Understanding Results ... 16
  - Benchmark Criteria ... 17
- More Information about IOmark-VM ... 17
- About Evaluator Group ... 17

Executive Summary

This document is the official benchmark report for the tested configuration using Datrium DVX 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 Datrium DVX storage with Datrium compute nodes running DVX software:

- A -node cluster achieved 8,000 IOmark-VM's at a cost of $667.01 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 d 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

Datrium DVX

DVX delivers the world's simplest and most powerful private cloud platform by converging compute, virtualization, primary and secondary storage in a radical new way. DVX splits IO processing from durable capacity using two basic blocks: compute nodes and data nodes. Add compute nodes for IO speed and data nodes for capacity to achieve independent, un-bottlenecked scaling. All data services such as deduplication, compression, erasure coding, snapshots, replication and encryption scale with compute nodes, are always-on and ultra-efficient for zero-knob data management at scale.

Datrium DVX provides the following enterprise features:

- Scale elastically from one compute node & data node to 128 compute nodes & 10 data nodes in a single DVX system.

- Use existing hosts & flash or buy turn-key Datrium compute nodes for single point of support.
- Keep compute nodes stateless and mix-workloads with back-up grade efficiency and economics.
- Snap, clone, restore or replicate any VM, vDisk, ISO file or container volumes alone or as self-updating policy group with all operations fully data efficient for extreme agility.
- Consolidate dev to prod IT across vSphere, RedHat, KVM and bare-metal docker containers.
- Fully WAN efficient, SSL encrypted Elastic Replication where all compute nodes transfer data from local flash using aggregated network bandwidth.
- Secure data with 100% SW-based encryption with built-in key manager. Data encrypted in-flight (from compute node to data node) and at-rest with full data reduction.
- Modern HTML5 GUI and a powerful vCenter plug-in (Flex and HTML5). Built-in real-time VM/container analytics & monitoring history. Integrated call home for proactive support.

IOmark-VM-HC Test Summary

For the tested configuration, the following data is provided.

Item	Value
Testing Identifier:	VM-HC-171024-b (revised version)
Tested Product(s):	10 @ Datrium DVX F12X2 storage nodes • 23 TB SSD raw, 8 x 10 GbE, 2 cntrlrs. 50 @ Servers, each with: • 2x Intel E5-2697v4 processors • 4x SSD cache + Datrium DVX software • 128GB RAM • VMware 6.0 U3 ESXi hypervisor
Priced Products:*	10 @ Datrium DVX F12X2 storage nodes • 23 TB SSD raw, 8 x 10 GbE, 2 cntrlrs. 60 @ Datrium Compute Nodes, each with: • 2x Intel Xeon Scalable 6148 processors • 4x SSD cache + Datrium DVX software • 512GB RAM VMware ESXi 6.0 hypervisor
Test Sponsor:	Datrium
Auditor:	Evaluator Group Inc.

Table 1: Test Identifier Information

* **Note:** For IOmark-VM-HC, 'tested' and 'priced' configurations may differ, as IOmark-VM tests VMmark storage workloads, but does not fully utilize the compute resources to the same level as the storage. In order to accurately report all resources required, the tested storage configuration must be matched with a VMmark workload to determine the server resources required. This is done to protect customers

who wish to use these configurations as guidelines for deployment. Thus, the priced configuration matches the tested storage workload with appropriately sized compute resources. Server resources are shown only for pricing purposes (total list price, $/VM). All other metrics shown are based on the tested configuration.

Item	Value
IOmark-VM Version:	Version: IOmark-VM 3.8
Testing Completed:	October 2017
Equipment Availability:	October 2017
Audit Certification Date:	24, October 2017
Report Date:	24, October 2017

Table 2: Test Revision and Dates

IOmark-VM-HC Results

Shown below are the IOmark-VM-HC results for the 'tested products' configuration. 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 and uses the 'priced products' configuration. Table 3 below shows an overview of the IOmark-VM-HC results.

Tested Product Results					Priced Product Results	
IOmark-VM Total VM's	**IOmark-VM Response Avg.**	**Available Capacity (wo/ data-reduction)**	**Used Capacity (w/ data-reduction)**	**Tested RAID Level**	**Total List Price**	**IOmark-VM-HC $ / VM**
8,000	2.11 ms	160 TB	125 TB	RAID 6	$5,336,112	$667.01

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 1,000 application sets yields 8,000 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:

Tested Product Results					
IOmark-VM Sets	**Read Resp. Average**	**Write Resp. Average**	**# vCenter Clone and Deploy**	**# vCenter storage vMotion**	**Hypervisor Workload Score (1 - inf.)**
1,000	0.76 ms	2.46 ms	6	5	1.5

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

Hardware Features	Value
Rack Footprint*	Tested: 50 Servers + 10 Datrium flash Data Nodes = 45 U (25 + 20) Priced: 60 Datrium Compute Nodes + 10 Datrium flash Data Nodes nodes = 80 U (60 + 20)
Number of Nodes per Appliance	N/A: Compute nodes and Datrium data nodes are independent
Number of Drives per Appliance	Cache: 4 Samsung PM863a / compute node = 200 total Capacity: 12 Samsung PM1633a / data node = 120 total
CPUs per system / total*	Tested: 2 x Intel E5v4 82 Ghz / node → 100 sockets; 4,140 Ghz Priced: 2 x Intel Xeon 96 Ghz / node → 120 sockets; 5,760 Ghz
Memory*	Tested: 128 GB / node → 6.4 TB for 50 nodes Priced: 512 GB / node → 30 TB for 60 nodes
Networking Ports (1/10 GbE)	Tested: 100 @ 10GbE Ports (2 Per Compute Node) Priced: 120 @ 10GbE Ports (2 Per Compute Node)

Table 5: Hyper-Converged Hardware Features

- **** Note:** The tested configuration is shown for transparency. However, for price and official configuration purposes, the "Reported" configuration required the CPU and memory resources as noted. The determination is made based on the referenced system running VMmark 2.5.*

The tested configuration utilized multiple servers with Datrium software and 10 Datrium F12X2 flash Data Nodes. The priced configuration consists of 60 Datrium Compute Nodes running Datrium software plus the tested 10 Datrium flash Data Nodes.

The CPU and memory necessary to support the VMmark 2.5 workload was determined to require 60 compute nodes, using 2 x Intel Xeon Scalable 6148 processors and 512 GB of DRAM. It is the determination of Evaluator Group that this configuration is capable of achieving 1,000 VMmark tiles based upon the published results of an Intel Xeon scalable 2 node system achieving 34 VMmark tiles, which translates into 1,020 tiles for a 60-node configuration. The referenced VMmark results are shown in the footnote below[1].

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

[1] VMmark 2.5 reported results: HPE 34 Tiles with Xeon Scalable processor

Hypervisor Configuration for IOmark-VM-HC Workload

- A single shared NFS mount point was accessed from all server nodes in the cluster
- Virtual disks were created for each of the reported IOmark-VM's certified (8,000 total)
- RAID level is not a tunable parameter and default Datrium policy is thin provisioned disks
- Dedupe, compression and encryption were enabled during testing

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 6.0 U3
Number of interfaces to the storage system:	Tested: 2 @ 10 GbE per compute node (100 total)
Connectivity to storage system:	4 @ 10Gb Ethernet / Datrium storage node
Hypervisor storage protocol used:	NFSv3
Hypervisor version:	VMware ESXi 6.0 U3
Thin provisioning:	Yes
Hypervisor Storage Access:	NFS datastore
Datastore Filesystem:	NFS access to Datrium DVX
VAAI:	N/A
SATP:	N/A
PSP:	N/A
Total capacity of system allocated to IOmark-VM:	10 DVX Nodes: 10 * 16 TB = 160 TB usable; 320 TB effective with 2:1 data reduction

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 NFS datastore was created using the pooled capacity across all 10 DVX storage nodes
- Each VM's VMDK was allocated using "thin provisioning" per Datrium 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	Datrium DVX 4.0.1.0
High Availability Access to all LUNs	Yes (active / passive)
Total raw capacity of system under test (SUT)	230 TB (across all nodes w/o data reduction)
Total usable capacity of system under test (SUT)	320 TB – 960 TB (2x – 6x data reduction rates)
NFS Mount Points	1 (Single shared NFS mount point for all nodes)
Thin provisioning:	Yes
RAID Level(s)	Erasure Coding 2 parity (DVX standard, no setting)
Total Cache Capacity:	Read cache per host, NVRAM write cache per DVX
Read Cache:	Tested: 384 TB (200 Samsung PM863a @ 1.92 TB) Priced: 460.8 TB (240 Samsung PM863a @ 1.92 TB)
Write Cache:	1 NVRAM per DVX data node
VAAI Features Enabled:	No - NFS access
- Block Zero	N/A
- Full Copy	N/A
- HW Locking	N/A
Automated tiering within the storage system:	N/A
Data Deduplication:	Enabled
Data Compression:	Enabled
Data Encryption:	Enabled
Storage system clones / writeable snapshots:	No, enabled not used
Storage Media Utilized:	HDD with NVRAM in DVX storage node
- SSD's (Note: Capacity Tier only, Cache noted above)	NVRAM cache, 160 TB usable on 10 DVX nodes
- HDD	-

Table 7: Storage System Configuration Parameters

Configuration Diagram

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

With Datrium DVX, a single NFS mount point is presented to all nodes in the cluster. Thus, only a single datastore is required for all virtual disks for VMs. The VMDK's for each VM's workload all were allocated from the same DVX NFS datastore capacity pool across 10 storage nodes.

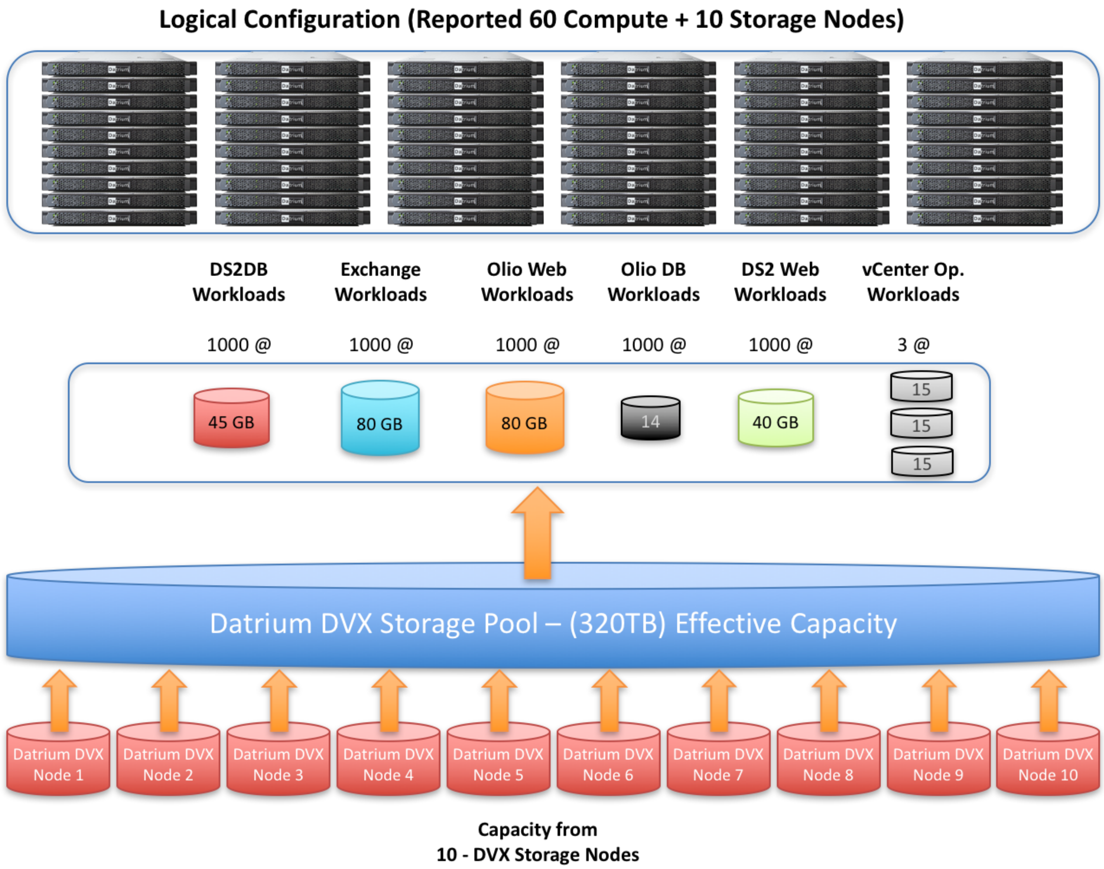

Figure 1: Logical System Configuration

Connectivity

The storage connectivity was two 10GbE ports per each ESXi host in the cluster. Testing did utilize a redundant HA configuration.

Note: *The priced configuration includes the cost of the 10 GbE ports within each server and 10 GbE ports on each DVX storage node. However, pricing does not include switches or other external network costs.*

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

Physical Configuration (Tested 50 Compute + 10 Datrium DVX Storage Nodes)

Figure 2: Physical System Connectivity

Priced Configuration at U.S. List Price

Item	Description	Qty.	Ext. List Price
1	Datrium Compute Nodes (2 @ 22 core 6184 CPUs, 512 GB DRAM + 2 @ 10 Gb Ethernet + 4 Samsung 1.92 TB SSD)	60 servers	$1,865,040.00
2	VMware vSphere Enterprise Plus	120 sockets	$419,400.00
4	Datrium DVX Host License	60 licenses	$720,000.00
5	Datrium DVX – F12X2 Data Node	10 nodes	$1,860,000.00
6	3 Years Support Datrium DVX data node HW	All Datrium	$169,704.00
7	3 Years VMware SW Support	120 sockets	$301,968.00
Total	**List Price HW + SW + 3 year service & support**		**$5,336,112.00**

Table 8: IOmark-VM-HC Price Information (60 Node Configuration)

Note: Support included for all hardware and software including storage 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.

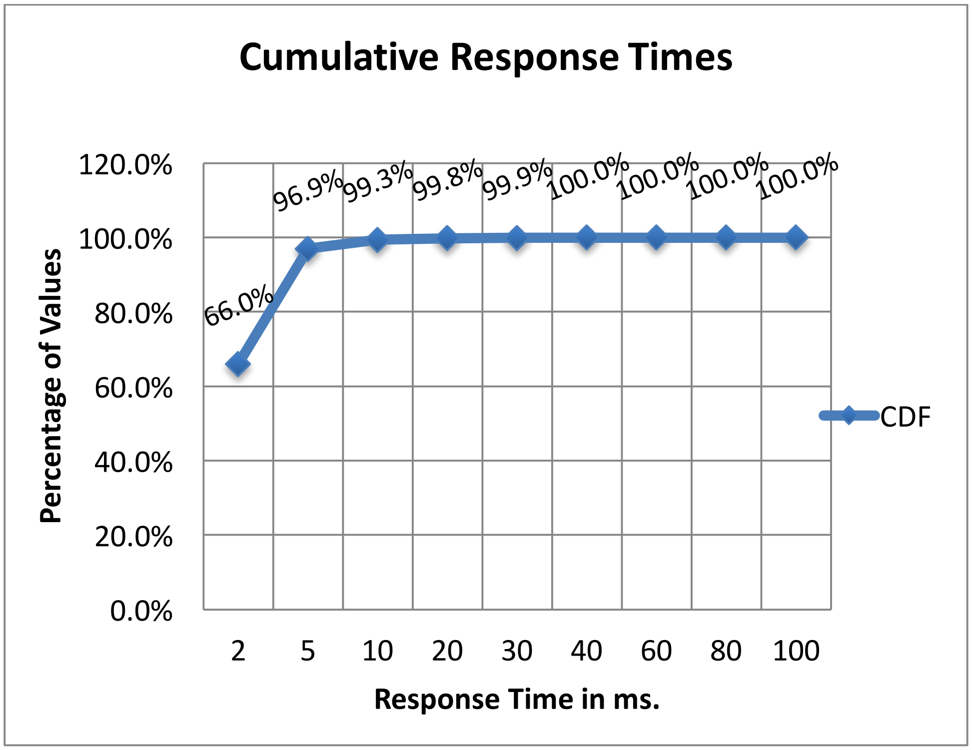

Figure 3: Percentage of Total Response Times at Measured Value

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

- Over 96% 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	1.06 ms	2.41 ms
Exchange Mail Server	1.03 ms	2.36 ms
Olio Web Server	0.88 ms	2.93 ms
Olio Database	0.84 ms	2.15 ms
DVD Store Web App 1	0.12 ms	2.14 ms
DVD Store Web App 2	0.12 ms	2.14 ms
DVD Store Web App 3	0.12 ms	2.14 ms
Windows Standby	0.12 ms	2.14 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.

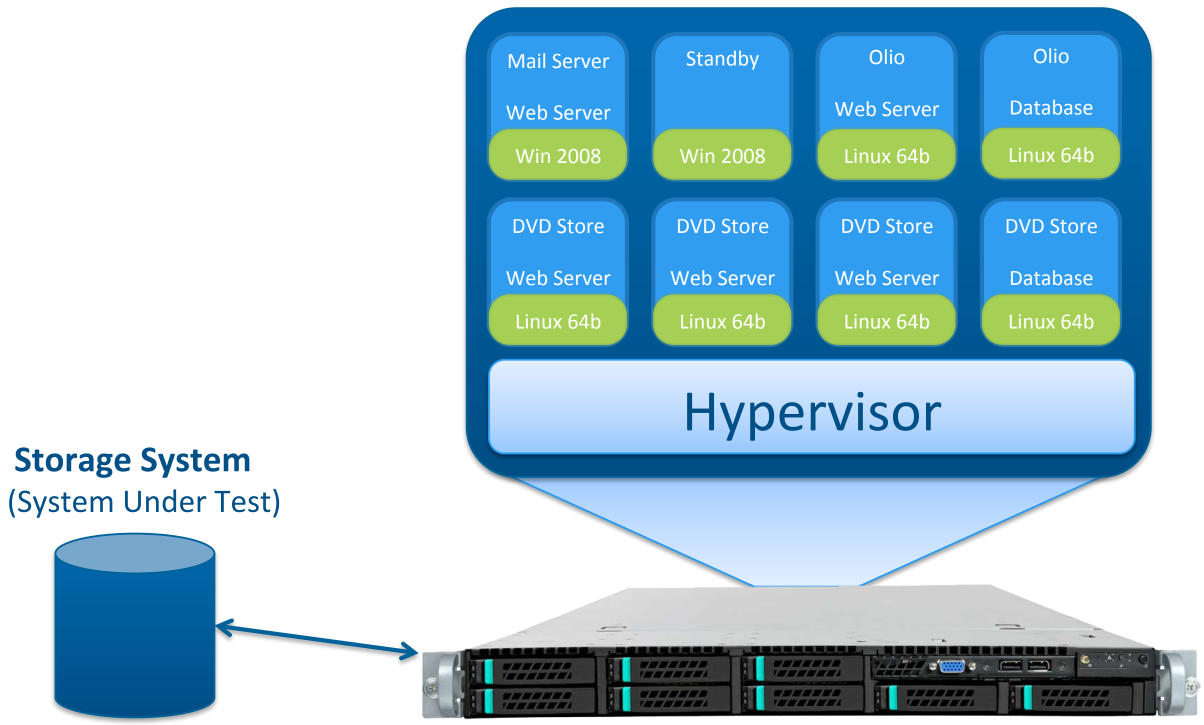

Figure 3: IOmark-VM Conceptual Overview

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 @evaluator_group