IOmark-VM

Tintri T5060 Storage System

Test Report: VM-160826-a

Test Report Date: 26, August 2016

Table of Contents

Executive Summary

This document is the official benchmark report for the tested configuration with a Tintri T5060 all-flash storage system running the IOmark-VM benchmark.

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

The results achieved by the Tintri system running IOmark-VM are as follows:

- The T5060 supported 480 IOmark-VM's at a cost of $468.75 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 system
  - 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 type must not exceed 20ms
  - 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

Vendor Supplied Product Description

Tintri simply stores virtualized workloads— offering a fully integrated VM-Aware Storage (VAS) system for virtualized enterprises and cloud. Balance your projects and workloads across high-performance All-Flash and award-winning Hybrid-Flash. Encrypt everything with a single click. Replicate individual VMs to remote locations. Sync child VMs with master VMs to speed development cycles. And manage it ALL from a single pane of glass. That's how you keep storage simple.

System design, including:

- Tintri VMstore UI, drill down to VM and virtual disk, in Tintri web-console or Hypervisor plugin
- Support for multiple, concurrent Hypervisors (VMware, Hyper-V, KVM and XenServer)
- Hypervisor-embedded storage for ease of management and deployment using existing tools
- Consistent performance and latency with all-flash T5000 series
- Operations managed per VM, including QoS, snapshots, clones and replication

Tintri all-flash T5000 series systems provide the following enterprise storage features:

- Capacity from 6 TB to 92 TB raw, and 17 - 308 TB effective capacity with data reduction
- Tintri all-flash arrays assign I/O on a per VM basis
- Quality of service (QoS) settings on individual VMs to guarantee performance
- Data reduction includes deduplication, compression, thin provisioning, and zero copy optimization

IOmark-VM Test Summary

For the tested configuration, the following data is provided

Item	Value
Testing Identifier:	VM-160826-a
Product(s):	Tintri T5060
Test Sponsor:	Tintri
Auditor:	Evaluator Group Inc.

Table 1: Test Identifier Information

Item	Value
IOmark-VM Version:	Version: IOmark-VM 3.8
Testing Completed:	July 2016
Equipment Availability:	August 2015
Audit Certification Date:	26, August 2016
Report Date:	26, August 2016

Table 2: Test Revision and Dates

IOmark-VM Results

Shown below are the IOmark-VM 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 results.

Config	IOmark-VM Total VM's	IOmark-VM Response Avg.	Tested Capacity	Tested RAID Level	Total Price	IOmark-VM $ / VM
T5060	480	4.33 ms	6.6 TiB (7.2 TB)	Tintri Default	$225,000.00	$468.75

Table 3: IOmark-VM Result Details

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 60 application sets yields 480 IOmark-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:

Config	IOmark-VM Sets	Read Resp. Average	Write Resp. Average	# vCenter Clone and Deploy	# vCenter storage vMotion	Hypervisor Workload Score (1 - inf.)
T5060	60	3.04 ms	4.44 ms	8	7	5.47

Table 4: IOmark-VM Passing Result Details

Tested Configuration Details

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

Hypervisor Configuration for IOmark-VM Workload

- A single Tintri T5060 system was used
- Multiple mount points were created to the two hosts used for testing
- A virtual disk was created for each of the reported IOmark-VM's certified (480)
- Data reduction was enabled, by generating data that is 2:1 compressible
    - RAID level utilized is not user selectable, used Tintri default settings

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

Storage System Parameter	Value
Hypervisor	VMware vSphere™ ESXi 6
Number of interfaces to the storage system:	4 Per Node (2 / controller = 4 total)
Connectivity to storage system:	2 @ 10Gb Ethernet / controller
Hypervisor storage protocol used:	NFSv3
Hypervisor version:	VMware ESXi 6.0U2 (3620759)
Thin provisioning:	Utilized in Tintri datastore
Hypervisor Storage Access:	NFS datastore
Datastore Filesystem:	NFS - filesystem access to Tintri
VAAI:	Yes, using Tintri VAAI plugin drivers
SATP:	N/A
PSP:	N/A
Total capacity of system allocated to IOmark-VM:	6.6 TiB usable (7.2 TB)

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 Tintri T5060 provided the pooled capacity across both of the host test nodes
- 6 datastores were created on each host for testing. (12 total)
- Each VM's VMDK was allocated using "thin provisioning" per Tintri and NFS default

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	Tintri 4.2.0.6
High Availability Access	Yes (active / standby)
Total raw capacity of system under test (SUT)	12 TB
Total usable capacity of system under test (SUT)	6.6 TiB (7.2 TB) usable, up to 35+ TB effective with data reduction
Datastores	Total of 6 datastores / host (No LUNs required)
Thin provisioning:	Yes
RAID Level(s)	Tintri default, RAID 6
Total Cache Capacity:	N/A
Read Cache:	N/A
Write Cache:	N/A
VAAI Features Enabled:	Yes
- NFS Full Clone	Yes
- NFS Extended Stats	Yes
- NFS Reserve Space	Yes
Automated tiering within the storage system:	N/A (T5060 is all-flash)
Deduplication or compression of data:	Yes, both in-line and always on
Storage system clones / writeable snapshots:	Yes, utilized storage clones during testing
Type of storage system clone:	Tintri native clones
Storage Media Utilized:	-
- SSD's	24 x 480 GB (includes spares)
- 15K RPM	NA
- 10K RPM	NA
- 7.2K RPM	NA

Table 7: Storage System Configuration Parameters

Configuration Diagram

The logical data layout of the test configuration is shown below in Figure 1. Since Tintri is an NFS datastore, block LUNs or volumes were not utilized. Instead, individual virtual disks were assigned to VM's as required, distributed across the 12 logical mount points. The VMDK's for each VM's workload were all allocated from the same Tintri Datastore capacity pool created by Tintri across the compute nodes.

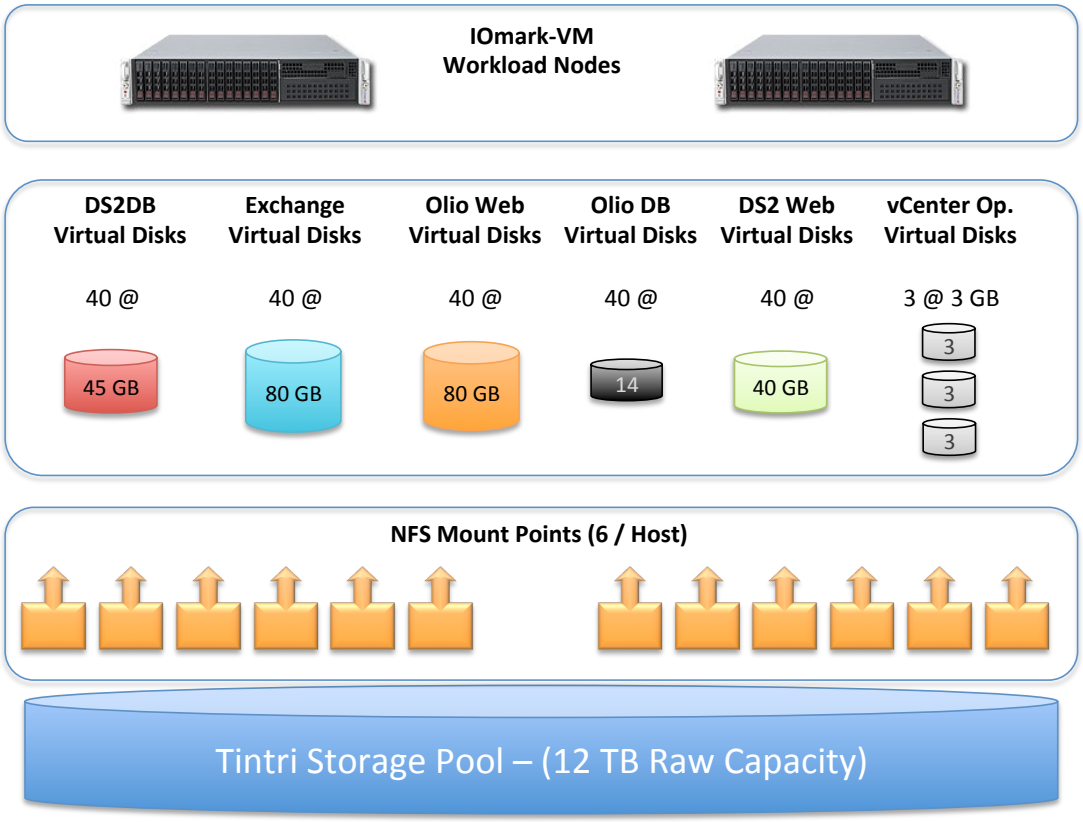

Figure 1: Logical System Configuration

Note: *Although a single NFS datastore could have been used, it was determined that utilizing more mount points improved performance. A total of 12 mount points were used, which equates to 12 host systems mounting a single Tintri shared datastore.*

Connectivity

Storage connectivity used was 10GbE, using LACP to bond the two network interfaces per controller. Each test host used 2 10GbE links to a 10 Gb Ethernet switch, for a total of 8 connections across both the two hosts and the dual controller T5060. Testing did not utilize a redundant HA configuration, although production deployment assumes connectivity to a HA network infrastructure.

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

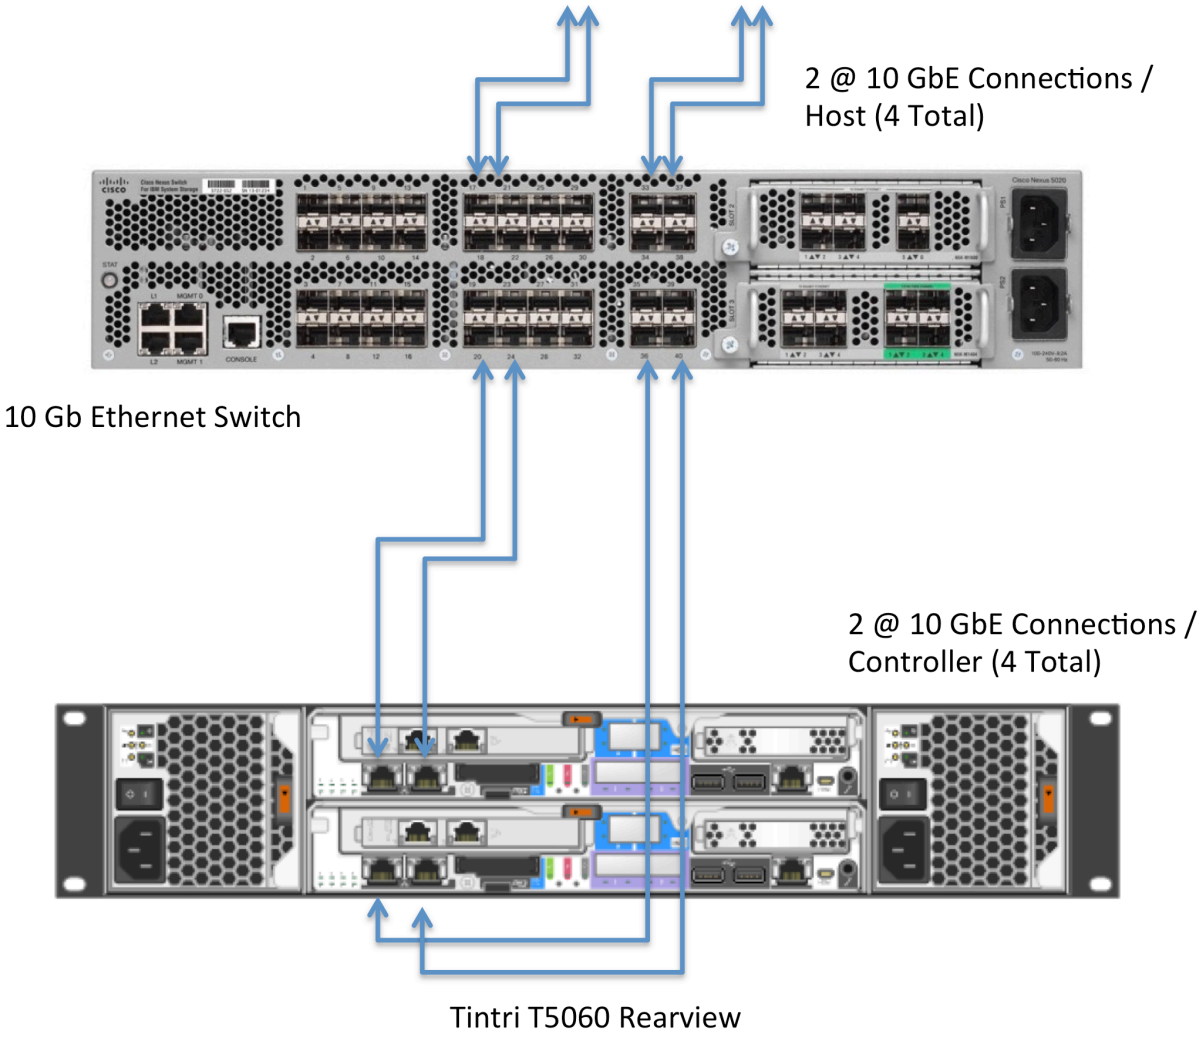

Figure 2: Physical System Connectivity

Tested Configuration Pricing

Item	Description	Qty.	Ext. List Price
1	Tintri T5060 (Includes base software)	1 system	$225,000.00
2	Tintri 3 years maintenance	Included	N/A
Total	**HW + SW + 3 year service & support**		**$225,000.00**

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

Note: Support included for all hardware and software

Detailed Results

IOmark-VM performance results are measured by application workload. The cumulative response times of all 480 applications that comprised the test workload are shown below in Figure 3, with average response times reported per application type below in Table 9.

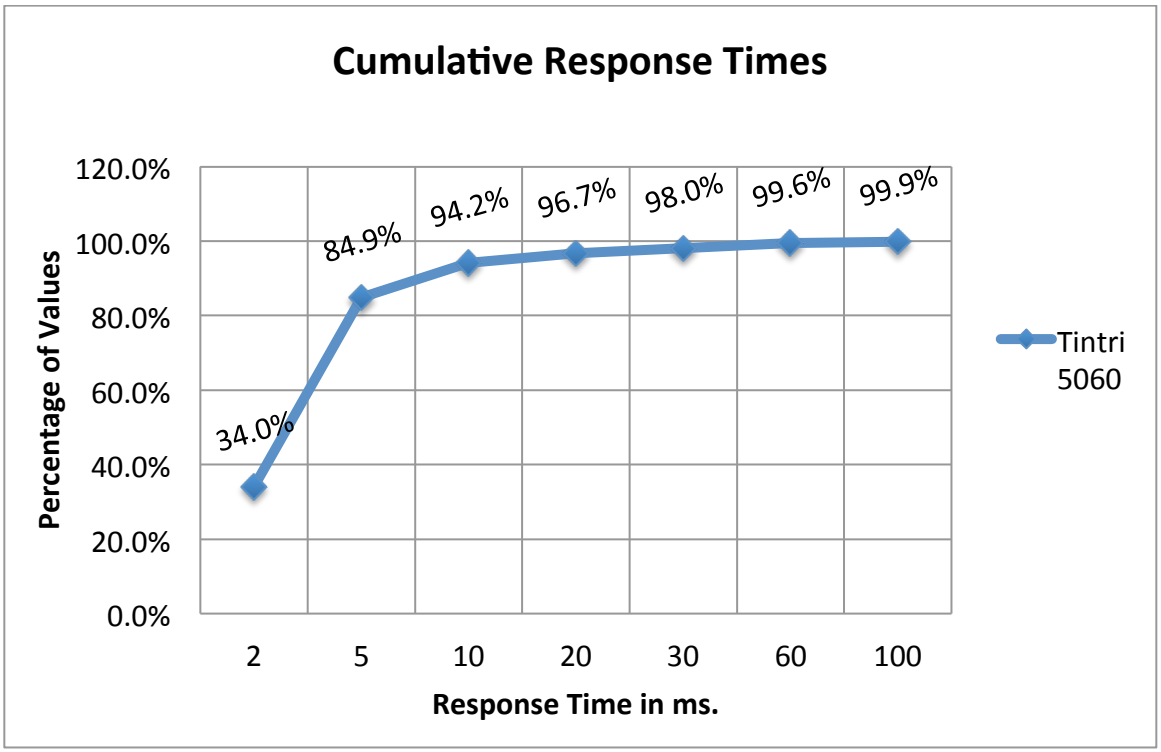

Figure 3: Percentage of Total Response Times at Measured Value

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

- Nearly 85% of response times were less than 5 ms.
- 96.7% of response times were less than 20 ms. (exceeding requirements)

Application Workload	Average Response Times
DVD Store DB	3.34 ms
Exchange Mail Server	4.27 ms
Olio Database Server	3.93 ms
Olio Web Server	5.38 ms
DVD Store Web App 1	4.57 ms
DVD Store Web App 2	4.57 ms
DVD Store Web App 3	4.57 ms
Windows Standby	4.57 ms

Table 9: Application Response Times by Workload Type

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.

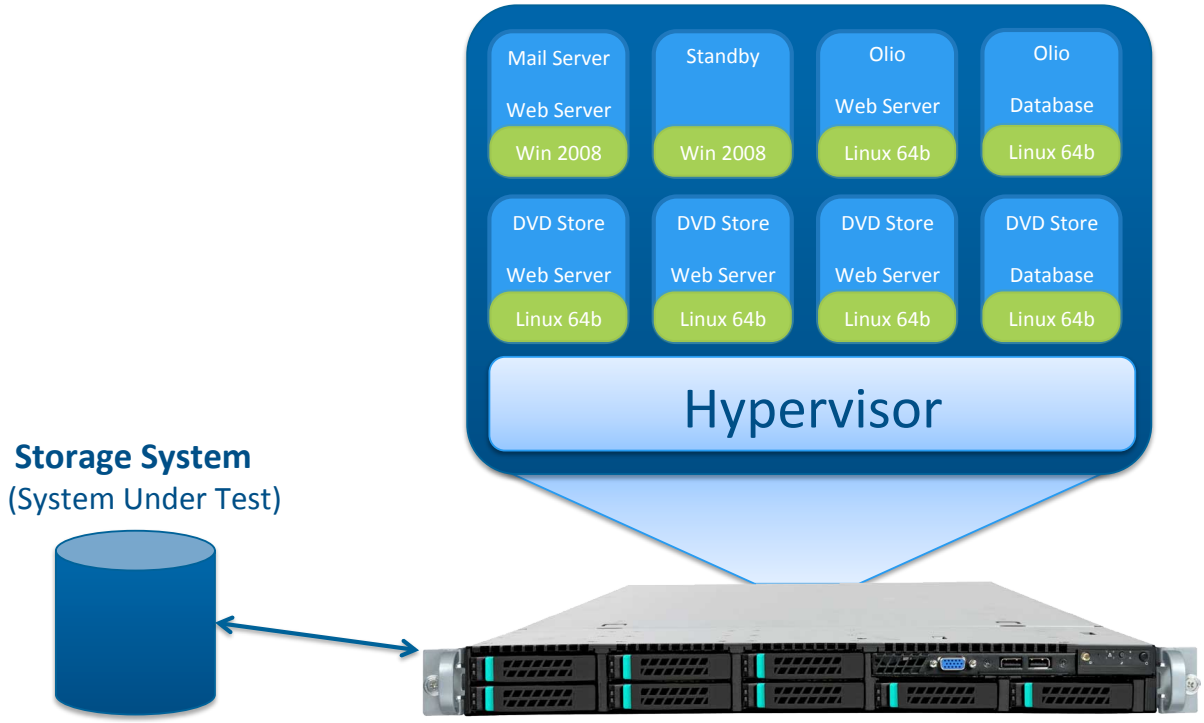

Figure 3: IOmark-VM Conceptual Overview

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 20ms
  - 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