# IOmark-VM


**Axellio FabricXpress**

Test Report: VM-190917-a

Test Report Date: 17, September 2019


# Table of Contents

# Executive Summary

This document is the official benchmark report for Axellio FabricXpress for Microsoft Azure Stack HCI, tested as a storage 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 storage solutions report IOmark-VM results and include compute and memory resources in addition to storage.

- Axellio FabricXpress achieved a record low price per IOmark-VM for storage systems
- IOmark-VM certifies FabricXpress cluster for 3200 VM applications at $126.10 / IOmark-VM
    - Configuration: 4 compute/storage nodes with 96 NVMe SSD drives
    - System total = $403,525.87 ($403,525.87 / 3200= $126.10 / 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

# Vendor Supplied Product Description

## Axellio FabricXpress

Axellio's FabricXpress (FX) solution is built around an extremely-dense, all-flash-based architecture optimized to run mission-critical workloads that now can be re-deployed away from inefficient, legacy data center infrastructure. Through a combination of unique hardware features and a strategic relationship with Microsoft, FX delivers the performance, financial, and operational benefits necessary to run scale-up and scale-out workloads in a single, simple HCI cluster architecture. Axellio's FabricXpress solutions are engineered specifically for high-performance requirements, but in a consolidated form factor that is easier to deploy, simpler to manage, and more cost efficient than earlier generations of HCI solutions.

Economic Efficiency:

- Less Space, Weight and Power (SWaP), drastically reducing datacenter costs
- Utilize enterprise software licensing structures with Microsoft to reduce $/VM costs
- Less OpEx through management simplicity, lower staff overhead, and easier scaling as you grow

Consolidation:

- Reduce physical and virtual footprint of your infrastructure
- Reduce server sprawl over traditional scale-out approaches
- Reduce storage latency and bandwidth bottlenecks

Performance:

- Handle unpredictable (and high-value) workloads like analytics and data warehousing
- High-speed PCIe architecture that enables dynamic scaling on a component-by-component basis
- Software-defined architecture that alleviates traditional performance bottlenecks in HCI-based servers, storage, and networks

## IOmark-VM Test Summary

For the tested configuration, the following data is provided.

| Item | Value |
|---|---|
| Testing Identifier: | VM-190917a |
| Product(s): | 4x Axellio FabricXpress for Microsoft Azure Stack HCI nodes |
| Test Sponsor: | Axellio |
| Auditor: | Evaluator Group Inc. |

**Table 1: Test Identifier Information**

| Item | Value |
|---|---|
| IOmark-VM Version: | Version: IOmark-VM 4.5.1 |
| Testing Completed: | June 2019 |
| Equipment Availability: | March 2019 |
| Audit Certification Date: | 19, August 2019 |
| Report Date: | 17, September 2019 |

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

| IOmark-VM Total VM's | IOmark-VM Response Avg. | Available Capacity | Used Capacity | Tested RAID Level | Total List Price | IOmark-VM $ / VM |
|---|---|---|---|---|---|---|
| 3200 | 1.35ms | 46 TB | 40 TB | 3 way mirror | $403,526 | $ 126.10/ VM |

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

***Note: Pricing shown is list price and does not include hypervisor licenses, pricing details in Table8.***

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 400 application sets yields 3200 VM's reported. The IOmark-VM workload may be used for IOmark-VM configurations.

The HyperV 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 HyperV 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 | # HyperV Clone and Deploy | # HyperV storage vMotion | Hypervisor Workload Score (1 - inf.) |
|---|---|---|---|---|---|
| 400 | 0.95 ms | 1.39 ms | 9 | 6 | 5.24 |

**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 S2D volume was created for the IOmark-VM's certified
- RAID level was assigned using 3-way mirror
- Axellio FabricXpress utilizes dynamic disks for each VM's virtual disk

Storage traffic utilized 100Gb/s Chelsio NIC, with VM traffic using the 2x 10Gb/s connections.

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

| Storage System Parameter | Value |
|---|---|
| Server version | Windows server 2019, Build 1809 |
| Number of interfaces to the storage system: | Storage traffic via 100GbE Chelsio,<br>VM traffic via 2x 10GbE |
| Connectivity to storage system: | 2x 10Gb VM, 2x Chelsio 100GbE for data |
| Hypervisor storage protocol used: | Storage Spaces Direct |
| Hypervisor version: | HyperV 10.0.17763.1 |
| Thin provisioning: | Yes - enabled |
| Hypervisor Storage Access: | Storage Spaces Direct |
| Datastore Filesystem: | CSVFS_ReFS |
| Total capacity of system allocated to IOmark-VM: | 96x 1.6TB SSDs, 153 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 S2D volume were created using the pooled capacity across all nodes
- The same S2D volume was utilized for administrative VMs and as clone and vMotion targets
- Each VM's VHDX was allocated using dynamic type for thin provisioning.

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 | Storage Spaces Direct with Windows Server 2019 |
| High Availability Access to all LUNs | Yes (active / active) |
| Total raw capacity of system under test (SUT) | 153TB (140TB Available Storage Pool Capacity before RAID) |
| Total usable capacity of system under test (SUT) | 140TB Raw storage available for all storage pools<br><br>Configured with 4x 11.5TB Volumes, 46TB Usable |
| Thin provisioning: | Yes |
| RAID Level(s) | 3-Way Mirror |
| Total Cache Capacity: | N/A |
| Read Cache: | N/A |
| Write Cache: | N/A |
| Automated tiering within the storage system: | N/A (All Flash) |
| Deduplication or compression of data: | Available, Not Utilized |
| Storage system clones / writeable snapshots: | Clones available, not used for testing |
| Type of storage system clone: | N/A |
| Storage Media Utilized: | NVMe SSD |
| - SSD's (Capacity Tier only, Cache noted above) | 96x Samsung PM1725a, 1.6TB, Dual-Ported NVMe |
| - HDD | N/A |

**Table 7: Storage System Configuration Parameters**

## Configuration Diagram

The logical data layout of the test configuration is shown below in Figure 1. A single storage pool was created, from which 4 identically sized volumes were created, one per host for the VM's virtual disks.

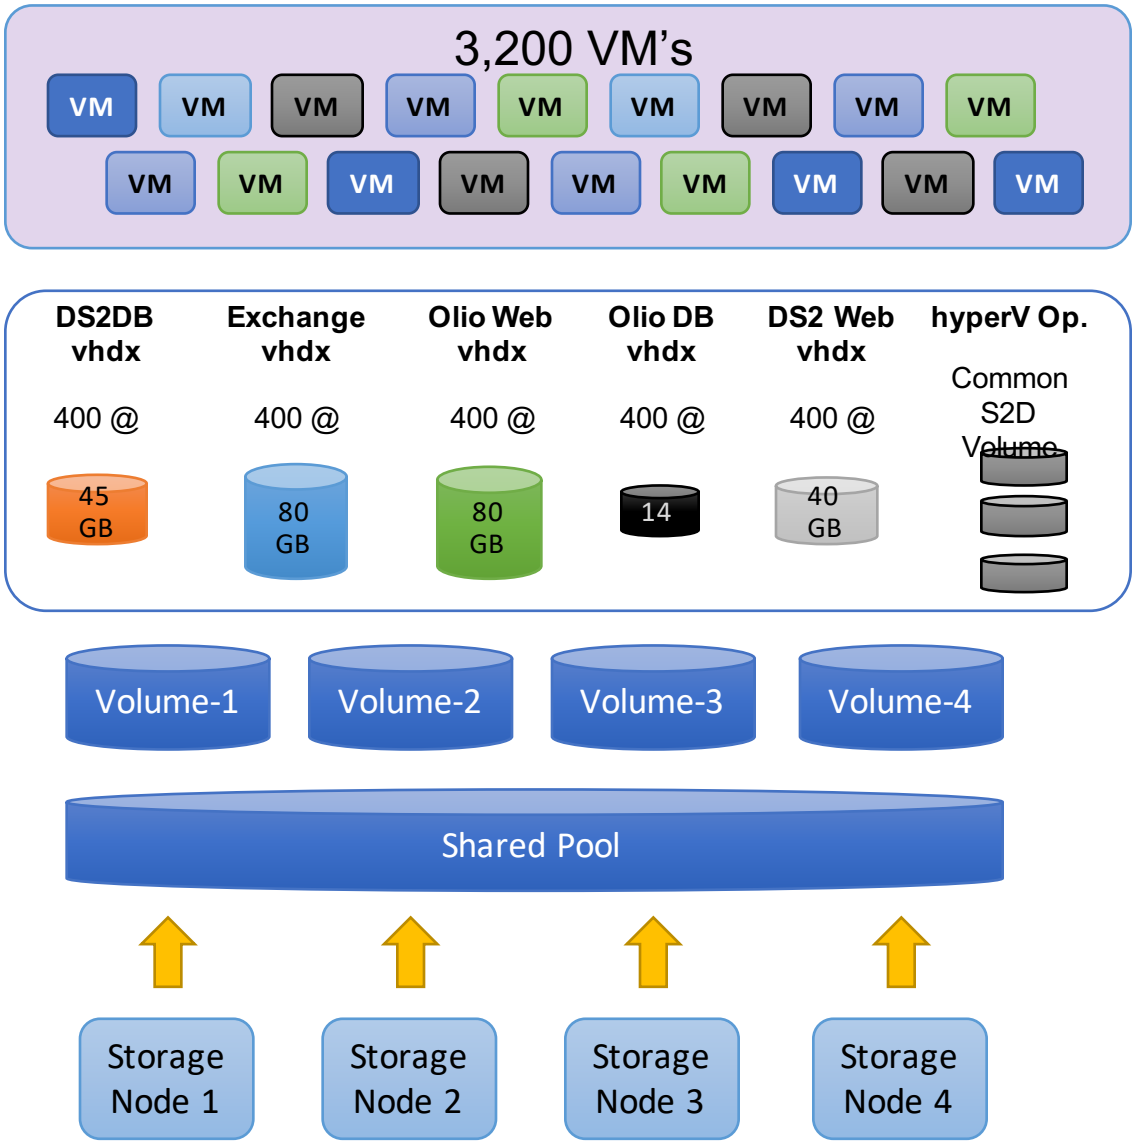


**Figure 1: Logical System Configuration**

## Connectivity

The storage connectivity was a 10Gb network for the "public" VM communications, and the 100Gb/s connections were for the "backend" network. The test setup utilized a Mellanox 100Gb/s switch for the backend network (fibre). Each node has a Chelsio T62100-LP-CR, dual ported 100Gb/s Ethernet card (offload enabled), and 2x 10Gb/s on-motherboard connections.

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

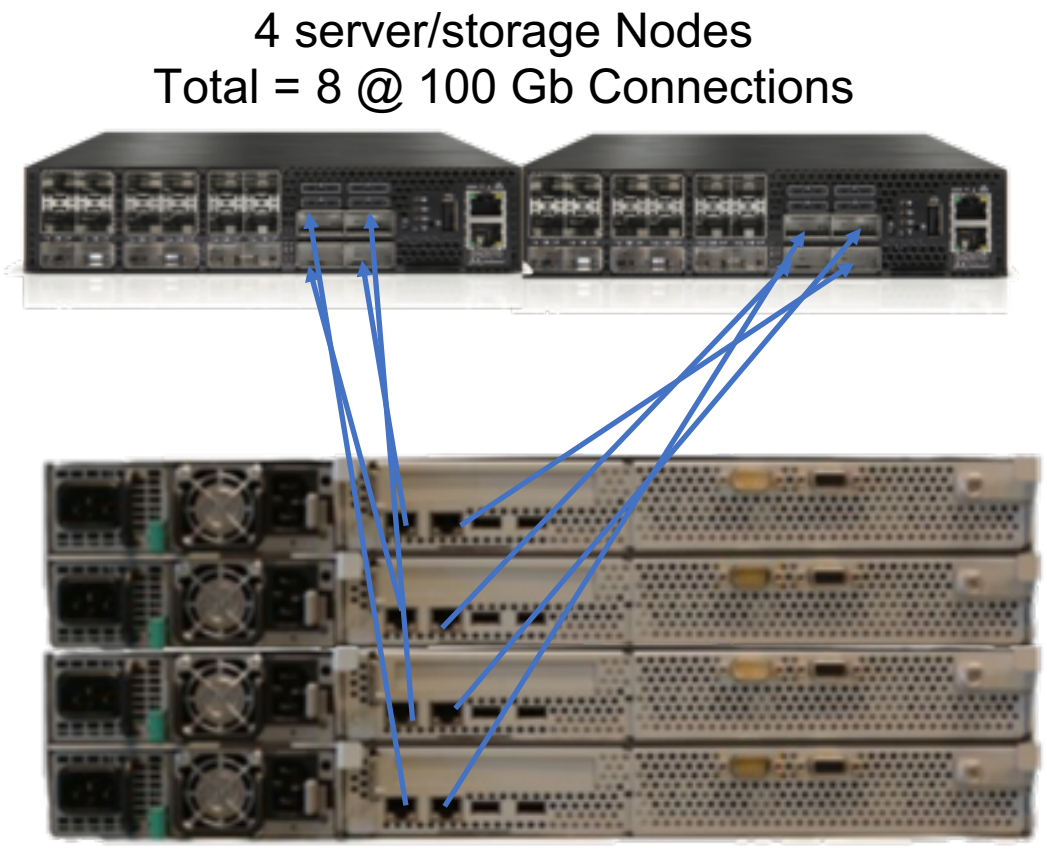


**Figure 2: Physical System Connectivity**

## Certified Configuration Pricing

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

| Item | Description | Qty. | Extended Price |
|---|---|---|---|
| 1 | Axellio FX-1000 Server Node (2 X E5-2699-v4 CPU @ 2.2 GHz, 512GB RAM @ 2666 Mhz, 2-port PCI add-in NIC) | 4 | $80,927.70 |
| 2 | Axellio FabricXpress SSD Carrier (8 X 1.6 TB NVMe drives) | 12 | $169,521.32 |
| 3 | Axellio FabricXpress FX-1000 Chassis | 2 | $43,258.14 |
| 4 | Axellio FX-1000 Advanced Support (NBD) – 1 year | 2 | $23,463.97 |
| 5 | Onsite Installation/day | 2 | $5,263.16 |
| 6 | Mellanox Spectrum MSN2010-CB2F 20-port switch | 2 | $12,631.58 |
| 7 | Windows Server 2019 Datacenter Licensing per 2 cores | 88 | $67,760.00 |
| 8 | Shipping & Handling | | $700.00 |
| **Total** | **List Price** | | $403,525.87 |

**Table 8: IOmark-VM Price Information (Pricing provided by Axellio)**

## 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.

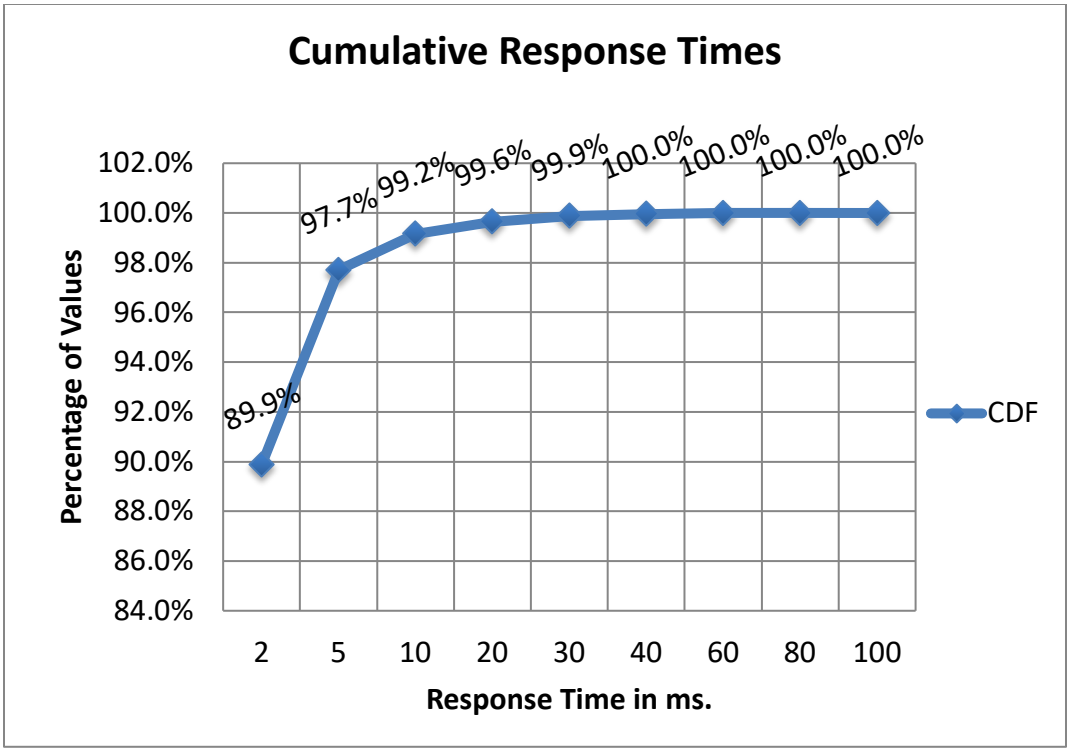


**Figure 3: Percentage of Total Response Times at Measured Value**

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

- Nearly 90% of response times were less than 2 ms for the cluster
- 97.7% of response times were less than 5 ms for the cluster
- Over 99% of response times were less than 10 ms for the cluster

| Application Workload | Avg. Response Time | |
|---|---|---|
| | Read | Write |
| DVD Store DB | 2.22 ms | 2.32 ms |
| Exchange Mail Server | 1.46 ms | 1.75 ms |
| Olio Web Server | 0.60 ms | 0.93 ms |
| Olio Database | 0.47 ms | 0.79 ms |
| DVD Store Web App 1 | 0.02 ms | 1.18 ms |
| DVD Store Web App 2 | 0.02 ms | 1.18 ms |
| DVD Store Web App 3 | 0.02 ms | 1.18 ms |
| Windows Standby | 0.02 ms | 1.18 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.

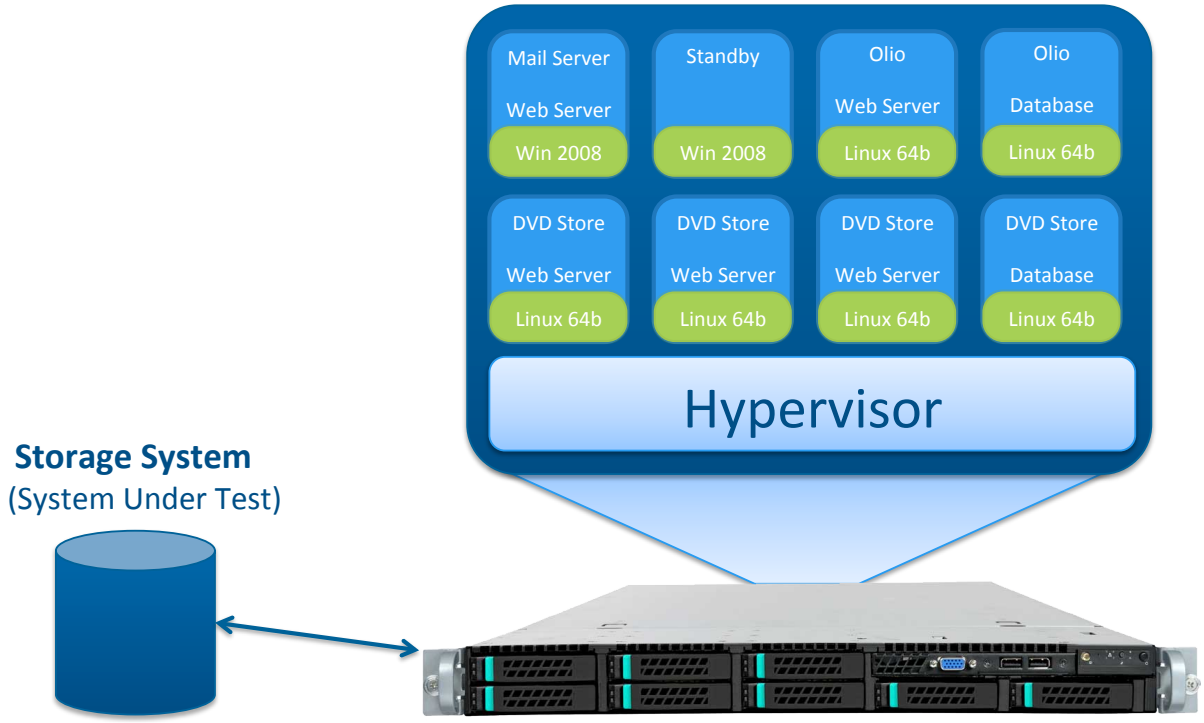


**Figure 4: IOmark-VM Conceptual Overview**

## 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 - HyperV required | 15 GB |
| **Hypervisor Storage Migration** | N/A - HyperV 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 HyperV 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*