EMC XTREMIO STORAGE ARRAY 4.0 AND VMWARE vSPHERE 6.0: SCALING MIXED-DATABASE WORKLOADS

Support growth with minimal cost

The EMC® XtremIO® array with always-on data services reduced

cost per IO by up to **43%**

cost per logical GB by up to **73%**

when creating three copies of our database workload.

The EMC XtremIO array with VSI for VMware® VSphere® supported consistent IOPS performance with each additional set of database VMs.

As organizations use more resources and face challenges managing database storage, administrators and storage teams using traditional storage arrays react by planning capacity, tuning performance, purchasing more storage space, and putting limits on usage. EMC offers a proactive alternative that can help you overcome these challenges: The all-flash XtremIO storage array with inline deduplication and compression data-reduction technologies.

Organizations that consolidate databases into the XtremIO get the following key benefits:

- Scalability for growth – database copies using the same data need less physical storage space than addressable (logical) space
- Performance boost – all-flash storage means greater computational speed
- More value as you scale – thanks to its use of addressable (logical) space rather than physical space, the more databases and performance the array can handle, the more it increases in value
- Ease the virtualized management burden – EMC offers integration tools such as EMC Virtual Storage Integrator (VSI) for a vSphere environment

In the Principled Technologies (PT) datacenter, we consolidated three different database applications into a single X-Brick® (the building block of XtremIO) supporting a virtual VMware® vSphere® 6.0 environment: Oracle® Database 12*c*, Microsoft® SQL Server®, and IBM® DB2®. We then measured capacity and IOPS performance as we scaled up the number of workloads and analyzed how storage space and performance improved the value of the array in terms of cost per logical GB and cost per IOPS.

SPACE SAVINGS & RELIABLE PERFORMANCE INCREASE THE VALUE OF THE EMC XTREMIO STORAGE ARRAY

In a chargeback scenario, an IT organization purchases an array to function as dedicated storage for various branches of the company. The IT department provisions storage to various departments within the organization, as requested, in a pay-for-what-you-consume service model. The multitenancy, chargeback service model ensures that no array capacity is wasted. In these scenarios, also known as "responsibility accounting," the XtremIO array can be especially valuable. Alternatively, if a business is operating in a showback environment, XtremIO's scalability can reduce IT costs for each tenant.

In the chargeback scenario, database administrators and other decision-makers in the chargeback scenario determine how much storage capacity the databases of the departments will need. The databases exist virtually on a reliable virtualization platform, such as VMware vSphere 6.0, and will often have the most demanding storage performance requirements, such as tier-1 storage (flash). The challenge is to find an array that can provide low latency with enough capacity to handle large workloads, potentially from different departments.

Enter the EMC XtremIO storage array. Designed with multi-controller architecture per brick, the XtremIO array features a dual-stage metadata engine and up to eight X-Bricks. Each X-Brick features two controllers, two battery backup units, and a 25-drive storage array enclosure for 2.5" drives. Each array is linearly scalable with clustering, using each additional X-Brick in lockstep. The XtremIO array is expandable while running, with no need for application outages.

According to EMC, "XtremIO transforms database workloads and storage management by reducing complexities and delivering the following benefits:

- All-flash performance that delivers high IOPS and sub-millisecond latency
- Consistent performance due to global data placement across all flash drives
- XtremIO Data Protection—the simple and efficient data protection mechanism—which features:
  - Lowest capacity overhead in the industry (only 8 percent)
  - Requires no hot spares and protects against dual SSD failures"[1]

[1] For more information about XtremIO, visit https://www.emc.com/en-us/storage/xtremio/benefits.htm.

In a chargeback scenario, the ultimate goal for the IT department is to recoup the capital expenditure (CapEx) of the XtremIO array while being able to take full advantage of its performance capabilities. This is easier to achieve with higher utilization of performance and physical capacity resources.

More tenants and more work for the array mean better utilization, more revenue for the IT organization, and faster recouping of CapEx. In addition, fully utilizing the storage resource drives down costs for all subscribers of the service and helps reduce the need to buy more storage.

When arrays have a large amount of available storage capacity, such as the two-brick XtremIO configuration, they can house many databases or other applications. In the chargeback scenario, departments in the organization are less likely to feel the need to go to third-party sources for additional storage needs. This helps the organization get a better value from the array and keeps the array generating revenue in the chargeback model scenario for IT.

INLINE DATA REDUCTION, THIN PROVISIONING, AND VSI

The XtremIO array can free up storage space through thin provisioning and inline data reduction. Thin provisioning means the system uses space only when necessary. It does not reserve empty blocks, so those blocks are available for other data, thereby increasing the addressable capacity of the array. The XtremIO array is also content-aware and can place blocks anywhere in the system while still handling random-access patterns. Randomly placed data is especially beneficial for SSDs because it lessens the occurrence of hotspots, or areas of frequent data access, which can reduce the lifespan of SSDs.

The data-reduction features employ “always-on” inline data deduplication and compression. The XtremIO array deduplicates by not storing redundant data segments entering the storage. Inline data compression reduces the amount of physical data that needs to be written to storage after the data is deduplicated. This feature further reduces the space used on the storage as well as the number of writes hitting the SSDs, so the array can store more data. Inline data compression can also potentially increase the lifespan of the SSDs. Figure 1 compares the all-flash X-Brick with other common storage environments.

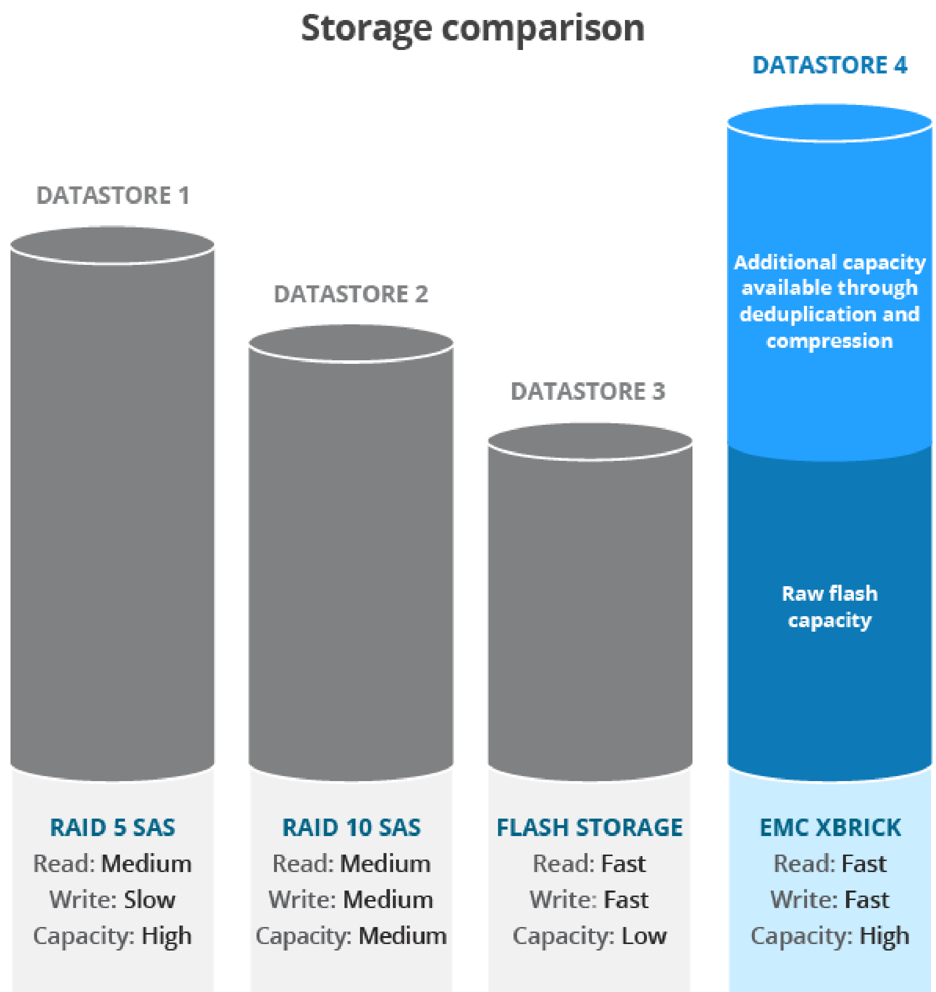

Figure 1: The EMC X-Brick can support fast reads and writes while keeping available storage space high (addressable capacity gained by deduplication and compression will vary depending on data).

XtremIO reports the deduplication, compression, thin provisioning, and overall efficiency of the entire array to the storage administrator through its GUI console. XtremIO also provides a Data Reduction Ratio that easily communicates to administrators the overall efficiency of the storage solution.

In the test environment, we used the EMC Virtual Storage Integrator (VSI) plug-in for vSphere to optimize host I/O settings. With the VSI plug-in, you can optimize your hosts, provision storage, clone VMs, and conduct other tasks when using EMC storage arrays that support the plug-in.

HOW WE TESTED

We used a single cluster of two X-Bricks to support three different, concurrently running, production-level database workloads:

1. DVD Store 2 (DS2), which simulates online transactions against a Microsoft SQL database
2. Silly Little Oracle® Benchmark (SLOB), which runs transactions against an Oracle 12*c* database
3. DB2_IOPS (DB2), a benchmark designed to stress IBM DB2 databases

We concurrently ran a single instance of each production-class database workload to create a performance baseline of IOPS and latency, and we used two SQL databases to generate a similar workload to the DB2 and Oracle database workloads.

We then used the XtremIO Virtual Copy (XVC) technology to copy the VMs and databases three times; measured total IOPS and physical storage footprint at each increase; and tuned each set of three VMs to a development level, running at a slightly lower IOPS count than the baseline production level.

Cloning VMs and databases on XtremIO storage invokes XVC. XVC technology allows copy operations to happen in-memory using metadata, reducing impact on the back-end. There are two main benefits of in-memory metadata copies:

- Instant copies of any application, including very large databases
- No initial capacity consumed

WHAT WE FOUND

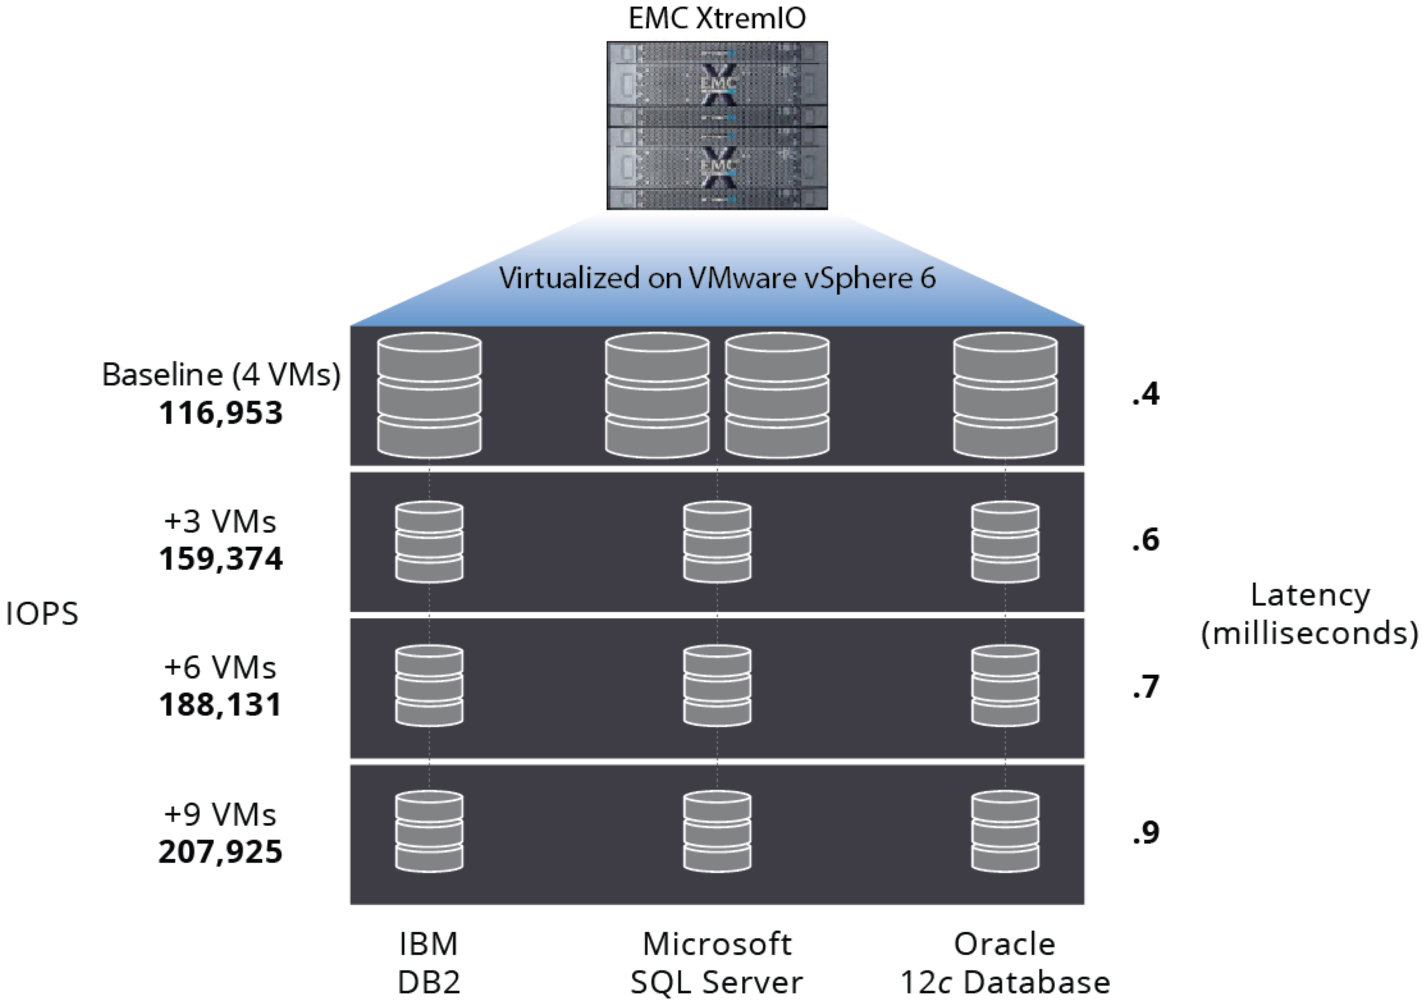

Figure 2: Testbed configuration overview with IOPS and latency results.

The all-flash storage of the XtremIO array hosted each database in our testing. The three production-level database workloads (baseline) provided 116,953 total IOPS with a 664GB physical storage footprint. Having this level of performance from the array helped deliver predictable throughput and low latency while running our workloads.

As we added the database VMs, the IOPS demand grew. Figure 2 shows the total IOPS for each phase of testing and the addressable architecture.

Maximizing storage space with the XtremIO array

Technologies such as inline deduplication, compression, and XVC optimize flash usage and minimize the amount of physical capacity used by databases and other applications. In each test phase, we monitored the physical and addressable capacity to validate XtremIO storage efficiency.

In Figure 3, physical capacity is the actual capacity used on the array and addressable capacity is the amount of capacity reported to the applications. XtremIO technologies such as inline deduplication and compression are transparent to applications, so the used physical capacity will frequently be smaller than the addressable capacity.

We found the least difference between the physical and addressable capacities in our baseline workload, because of how inline compression and deduplication work. The four databases in the baseline were unique, which minimized the impact of deduplication. In later workload scenarios, which used more virtual copies of each database, the difference in space increased because there was more duplicate data to address.

	Physical (GB)	**Addressable (GB)**	**Difference (Space savings - GB)**
Baseline	664	1,925	1,261
Baseline + three VMs	1,209	4,147	2,938
Baseline + six VMs	1,268	5,570	4,302
Baseline + nine VMs	1,595	7,213	5,618

Figure 3: As we increased the workload, the physical space increased only slightly as the addressable space increased.

Figure 3 shows a trend in which addressable capacity growth significantly outpaced physical capacity growth. We addressed the difference between physical and addressable capacity in GB of savings. In our tests with the baseline workload plus nine VMs, the space savings was 5,618 GB. The more virtual copies of databases and other applications you add, the greater the capacity savings can be.

Looking at the data on a more granular level shows the detailed space-saving benefits from scaling. XtremIO's built-in deduplication and compression can reduce the storage space required per VM or workload instance. In our testing, the storage space used per workload instance decreased steadily as we added VMs (see Figure 4). Although the physical footprint increased slightly as we increased the number of workloads, the physical storage space per workload instance decreased by 39 percent. This shows that despite a growing mix of databases, using the data-reduction technologies of XtremIO kept the amount of necessary data from rising drastically.

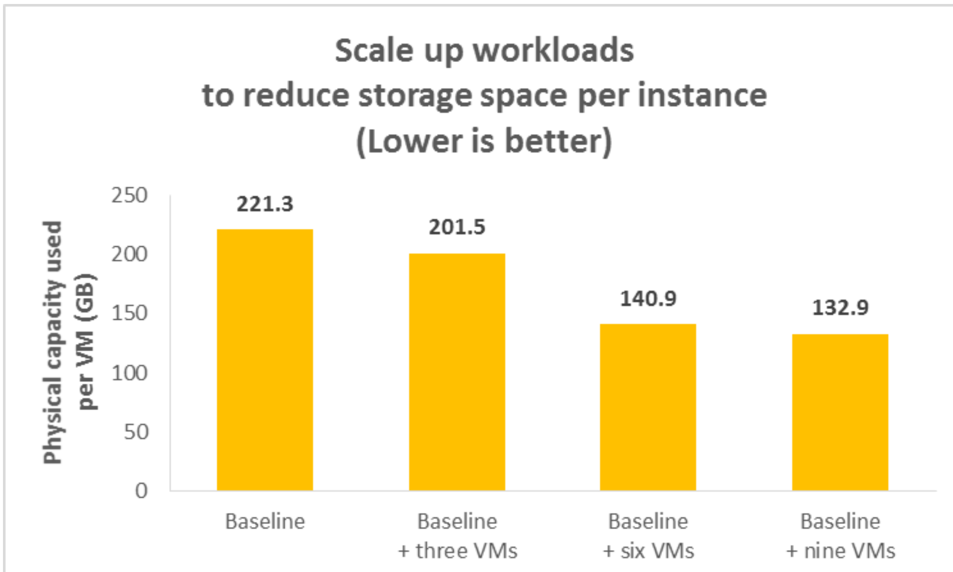

Figure 4: With each new set of VMs, the physical storage space used per workload instance decreased.

Consistent IOPS and latency

Production workloads normally have dedicated storage to help administrators ensure predictable, consistent performance. We wanted to see if this all-flash XtremIO array could drive the IOPS that the workloads demand while delivering sub-millisecond performance. In the first phase, we set up our baseline with the following production-level database workloads:

- Two Microsoft SQL Server databases targeting 40,000 IOPS combined
- One DB2 database targeting 40,000 IOPS
- One Oracle 12*c* database targeting 40,000 IOPS

With the three production workloads generating 116,952 IOPS, XtremIO delivered an average latency of .4 milliseconds. This baseline phase of testing shows a two-X-Brick XtremIO configuration can support a consolidated production workload across three different databases at performance levels expected of all-flash arrays. IT organizations looking to consolidate production-level database workloads can use XtremIO to deliver necessary IOPS and latency in a 13U footprint.

We then used XtremIO Virtual Copy Services to create instant copies of each database for the next three phases. Intending for these copies to simulate databases for test workloads, we tuned each set of the three test databases to generate fewer IOPS than the production-level databases, which more accurately simulates non-production databases. Each of the next three development-level phases of testing featured the addition of the following:

- One Microsoft SQL Server database targeting 15,000 IOPS
- One DB2 database targeting 15,000 IOPS
- One Oracle 12*c* database targeting 15,000 IOPS

With each addition of development-level VMs, the total IOPS increased while the average latency remained under a millisecond. This shows that the XtremIO array was able to meet IOPS demand while keeping response times low (see Figure 2 for latency numbers at each level of testing).

In terms of consolidation, XtremIO can manage multiple types of databases and workloads, in both test and production environments, without sacrificing performance. Many IT organizations have different service level agreements (SLAs), but for the demand of sub-millisecond performance in a densely consolidated storage platform, the XtremIO array can deliver.

The value of XtremIO

Scaling up the number of database workloads on an XtremIO array can get you a better value, which is desirable in a chargeback model as you add tenants. Figure 5 shows how adding workloads and using more addressable space can improve the value of XtremIO by reducing the cost per GB. As workloads use more addressable space and XVC technology repurposes copies of production, the cost per logical GB decreases. Looking at the data another way, the cost per logical GB at baseline was 3.74 times the cost of 13 databases.

We used the addressable space to determine the cost per GB.

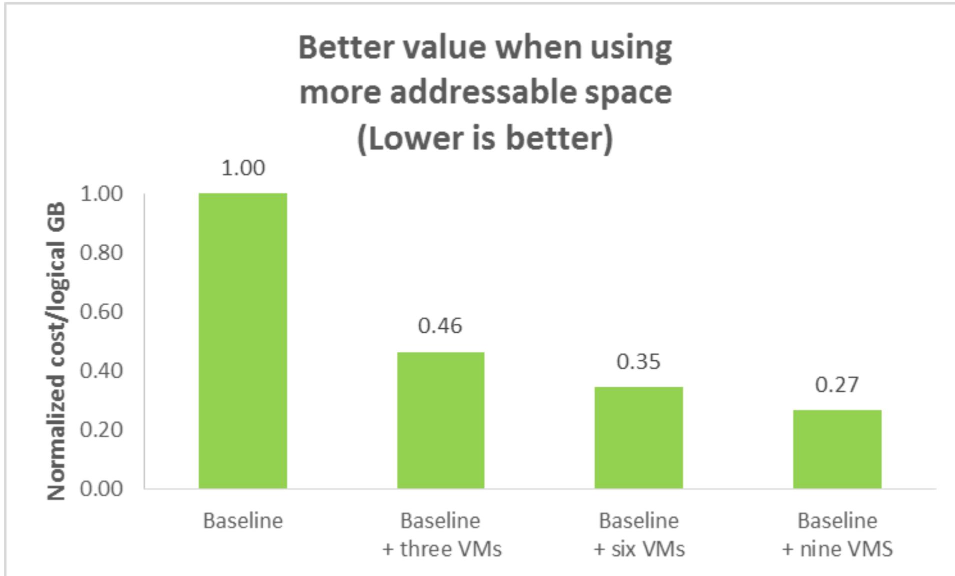

Figure 5: The normalized cost per logical GB for 13 database workloads was 73.3 percent lower than the normalized cost at baseline.

Getting greater performance in XtremIO decreased the cost per I/O, which would also be important in our chargeback scenario. Figure 6 shows that, like the cost per logical GB, the normalized cost per IOPS continued to drop as the workload increased. The increased IOPS from 13 database workloads lead to the greatest difference from baseline—43.8 percent lower cost per IOPS. Looking at the data another way, the cost per IOPS at baseline was 1.77 times the cost of 13 databases.

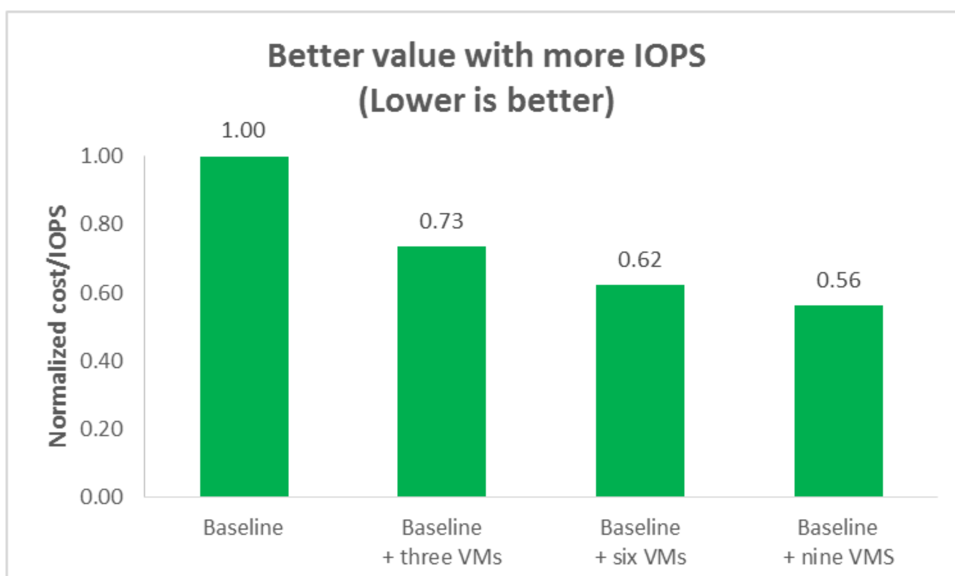

Figure 6: Based on our testing, the cost per IOPS of the array decreased as IOPS increased.

We used 1.595 TB of the array's 30.489TB physical capacity with the max level of 13 databases. There was available capacity to add more workloads, which could then have continued to decrease the cost per logical GB and the cost per IOPS.

CONCLUSION

In a mixed-application workload running in a cost-conscious environment, such as our chargeback scenario, IT organizations must be able to meet performance SLAs and consolidate applications. Database VMs need plenty of storage capacity and performance to handle the increased workload demands users place on them. Measures such as purchasing new arrays to meet these demands can be costly.

Thanks to its data reduction technologies, the EMC XtremIO storage array 4.0 saved storage space while supporting additional development-level database VMs in a VMware vSphere 6.0 environment. XtremIO substantially saved capacity by leveraging inline compression, inline deduplication, and virtual copies. The addressable capacity in our largest test run was 7,213 GB but took only 1,565 GB of physical space on the all-flash array. At our most I/O-intensive and largest-scale performance level, our 13 workloads generated 207,927 IOPS at an average latency of .9 milliseconds. Although we focused on increasing IOPS, latency remained under one millisecond in all of our mixed-application workload tests.

Based on our findings, scaling workloads, saving storage capacity, and delivering speedy all-flash performance can improve the value of the array. In the small capacity footprint at the 13-database level, the cost per addressable GB shrunk by 73 percent. Had our tests been larger and used more XtremIO capacity, we could have potentially found greater reduction in terms of price per GB. We also calculated the cost per IOPS and saw a 43 percent reduction at the 13-database level from baseline.

APPENDIX A – TESTING CONFIGURATIONS AND BENCHMARKS

About the EMC XtremIO 100% Flash Scale-Out Enterprise storage array

XtremIO is an all-flash scale-out enterprise storage array designed to substantially boost I/O performance. According to EMC, "XtremIO delivers new levels of real-world performance, administrative ease, and advanced data services for applications." Figure 7 shows the XtremIO user interface.

According to EMC specifications, the XtremIO 5TB, 10TB, and 20TB X-Bricks are populated by either 400GB or 800GB eMLC solid-state drives (SSDs). The Starter X-Brick has 13 400GB SSDs for a capacity of 5.2 TB and can be dynamically expanded online up to 10 TB. The 10TB and 20TB X-Bricks have 25 SSDs per system with capacities of 400 GB and 800 GB, respectively. With the current qualified six-X-Brick clusters orderable today, that allows for up to 120 TB of physical capacity.[2]

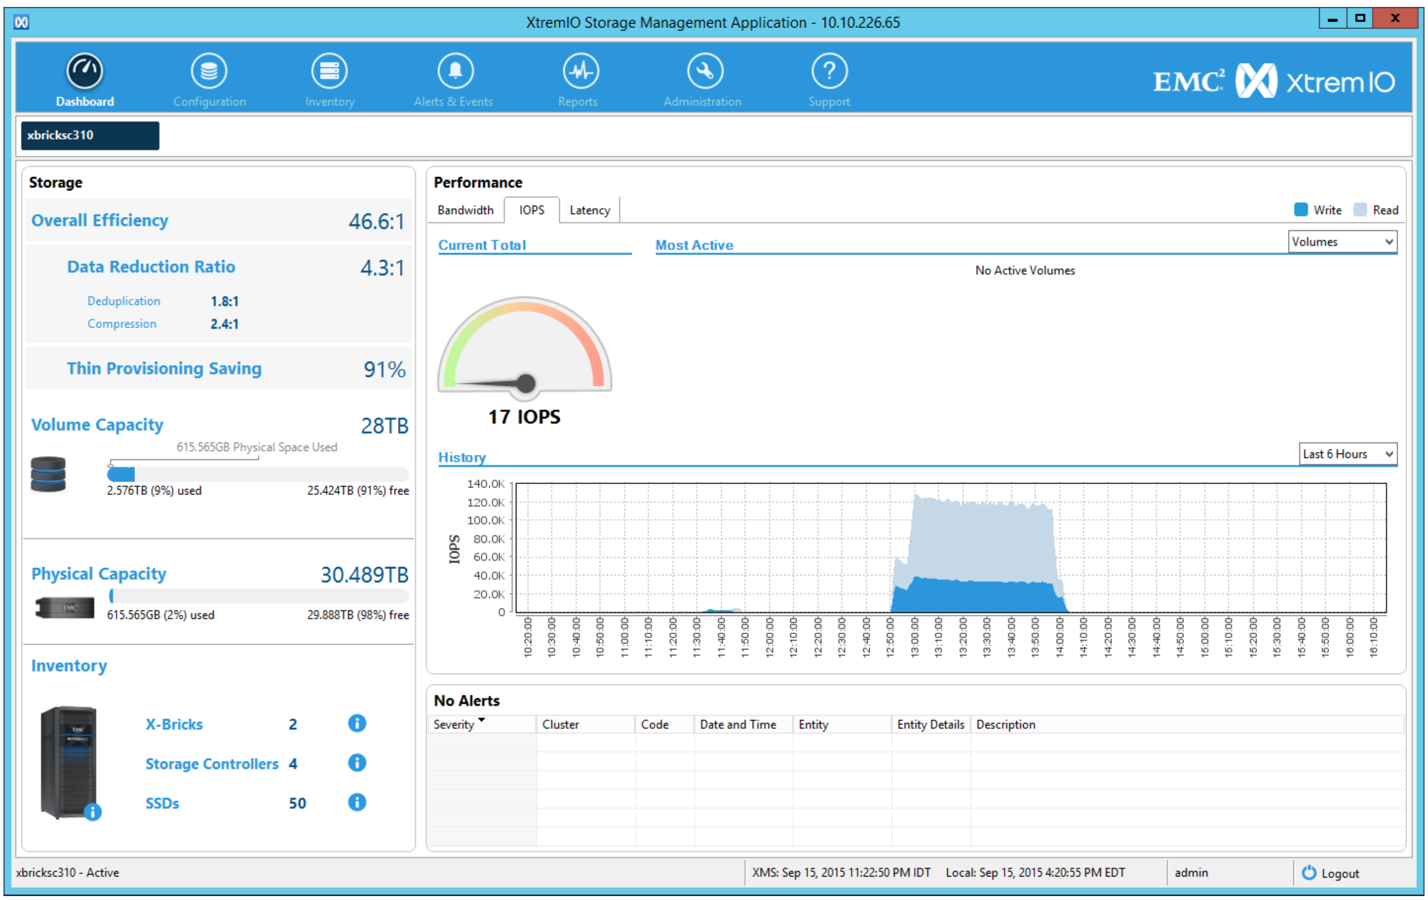

Figure 7: XtremIO Storage Management Application.

For features, XtremIO includes inline data deduplication,[3] compression, and XtremIO Virtual Copy (XVC) technology as part of its standard offering. Inline data deduplication and compression can potentially translate into

[2] xtremio.com/capacity
[3] xtremio.com/data-reduction

significant space savings across databases and applications. XVC technology was designed by EMC to be very space-efficient and to enable the instant availability of copies of a database.[4]

Business-critical databases and applications typically require a high degree of availability and integration with virtualized solutions. XtremIO's N-way active-active controller architecture is created to be extremely reliable. According to EMC, it has "no single point of failure."[5] XtremIO Data Protection (XDP) uses N+2 protection on every X-Brick, and seeks to provide protection against multiple SSD failures.[6] XtremIO integrates with VMware by providing a Virtual Storage Integrator (VSI) and VAAI plug-ins at no charge.

EMC protection technologies such as EMC RecoverPoint®, VPLEX®, or Data Domain® can integrate with XtremIO and extend the availability of your data. XtremIO's strong performance, space-saving technologies, built-in data protection, and integration with other EMC protection technologies can give your business a powerful option for virtualization storage. To learn more, visit www.emc.com/XtremIO, and for the latest technical papers, visit www.emc.com/everythingoracle.

About VMware vSphere

VMware vSphere is the flagship platform of virtualized hardware from VMware. vSphere allows companies to virtualize their server, storage, and networking resources, achieving a consolidation ratio in some cases greater than 15:1. Features included in vSphere include automated management and dynamic resource allocation. To learn more about VMware vSphere, visit www.vmware.com/products/vsphere/overview.html.

According to VMware, vCenter Server™ "provides a centralized platform for managing VMware vSphere environments." vCenter™ offers visibility into a vSphere infrastructure to analyze and fix issues. Features such as automated load balancing and out-of-the-box automation workflows create automated proactive management. vCenter can work with third-party solutions to extend its virtualization capabilities. For more information, visit www.vmware.com/products/vcenter-server/.

About our testing configuration

About our hardware configuration

Our testbed hardware configuration consisted of two XtremIO X-Bricks and eight bare-metal hosts. Each host had two Intel® Xeon® E5-2690 v2 processors and 256GB memory. Each host used a dual-port Emulex® LightPulse® LPe16000 PCIe® Fibre Channel adapter to connect to the EMC XtremIO X-Bricks with two Cisco® MDS 9148 fabric switches in between. We also deployed a single Brocade VDX 6740 10Gbps switch for management traffic.

[4] xtremio.com/snapshots_v2
[5] xtremio.com/high-availability
[6] xtremio.com/data-protection-2

About our storage configuration

Figure 8 shows the volumes on the storage that hosted our VM configuration files and data disks.

Volume	Size (TB)
OS1	8
DATA1	4
DATA2	4
DATA3	4
DATA4	4
LOG1	4

Figure 8: Volume information from our testing.

For each workload VM, we placed OS, data, and log VMDKs on various volumes. Figure 9 shows the vCPUs, memory allocation, and VMDK placement for each type of workload.

VM	vCPUs	Memory (GB)	OS1 (GB)	DATA1 (GB)	DATA2 (GB)	DATA3 (GB)	DATA4 (GB)	LOG1 (GB)
DVD Store 2	32	4	40	300				250
SLOB	8	4	50	50	50	50	50	100
DB2_IOPS	8	4	50	100	100			100

Figure 9: Specification of our test workloads.

About our workload configuration

For our DVD Store 2 workload VMs, we installed Windows Server® 2012 R2 and Microsoft SQL Server® 2014. We divided the production-level workload between two identical VMs in order to get more total IOPS from this workload. We used controller VMs hosted outside of the testbed to run the workload and collect performance counters. During the cloning phases, we only cloned one DVD Store 2 VM, as only one was necessary to create the desired development-level workload profile.

For our SLOB workload VMs, we installed Oracle Enterprise Linux® 6.5 onto one VM and installed Oracle Database 12*c*. We then installed the SLOB benchmark locally on the VM.

For the DB2_IOPS workload VMs, we installed Red Hat® Enterprise Linux 6.6 and the benchmark software, which includes an embedded version of the DB2 database.

About the DVD Store Version 2.1 benchmark

DVD Store 2 (DS2) models an online DVD store, where customers log in, search for movies, and make purchases. DS2 reports these actions in orders per minute that the system could handle to show what kind of performance you could expect for your customers. The DS2 workload also performs other actions, such as adding new customers, to exercise the wide range of database functions you would need to run your ecommerce environment.

For more details about the DS2 tool, see www.delltechcenter.com/page/DVD+Store.

About the SLOB 2.2 benchmark

The Silly Little Oracle Benchmark (SLOB) can assess Oracle random physical I/O capability on a given platform in preparation for potential OLTP/ERP-style workloads by measuring IOPS capacity. The benchmark helps evaluate performance of server hardware and software, storage system hardware and firmware, and storage networking hardware and firmware.

SLOB contains simple PL/SQL and offers the ability to test the following:

1. Oracle logical read (SGA buffer gets) scaling
2. Physical random single-block reads (db file sequential read)
3. Random single block writes (DBWR flushing capacity)
4. Extreme REDO logging I/O

SLOB is free of application contention, but it is an SGA-intensive benchmark. According to SLOB's creator, Kevin Closson, SLOB can also offer more than testing IOPS capability – actions such as studying host characteristics via NUMA and processor threading. For more information on SLOB, links to information on version 2.2, and links to download the benchmark, visit kevinclosson.net/2012/02/06/introducing-slob-the-silly-little-oracle-benchmark/.

About the DB2_IOPS benchmark

DB2_IOPS is a synthetic benchmark utility designed to stress DB2 databases. DB2 is an IBM database product designed to store, analyze, and retrieve data efficiently. DB2 is used across many different business applications to cover relational database needs.

APPENDIX B – SYSTEM CONFIGURATION INFORMATION

Figure 10 provides detailed configuration information for the XtremIO storage array.

Storage array	EMC XtremIO 4.0
Number of storage shelves	2 × 25 disks enclosure
Number of storage controllers	4 (two per storage shelf)
Controller OS	XtremIO OS 4.0.0-59_ndu
Disk model number	25 × Hitachi® HUSMM818
Disk size (GB)	800
Disk type	SAS SSD

Figure 10: Detailed configuration information for the storage array.

Figure 11 provides detailed configuration information for the test systems.

System	Intel Xeon processor-powered servers
General	
Number of processor packages	2
Number of cores per processor	10
Number of hardware threads per core	2
System power management policy	Default
CPU	
Vendor	Intel
Name	Xeon
Model number	E5-2690 v2
Socket type	FCLGA2011
Core frequency (GHz)	3
Bus frequency	8 GT/s
L1 cache	32 KB + 32 KB (per core)
L2 cache	256 KB (per core)
L3 cache	25 MB
Platform	
Vendor and model number	Intel S2600GZ
Motherboard model number	G11481-354
BIOS name and version	Intel SE5C600.86B.02.03.0003.041920141333
BIOS settings	Default
Memory modules	
Total RAM in system (GB)	256
Vendor and model number	Micron® 36JSF2G72PZ-1G9E3
Type	PC3-14900
Speed (MHz)	1,866
Speed running in the system (MHz)	1,866
Size (GB)	16
Number of RAM module(s)	16
Chip organization	Double-sided
Rank	Dual

System	Intel Xeon processor-powered servers
Hypervisor	
Name	VMware ESXi™ 6.0.0
Build number	2809209
Language	English
Fibre Channel adapters	
Vendor and model number	Emulex LightPulse LPe1600
Type	Dual-port 16Gb Fibre Channel

Figure 11: Configuration information for the systems used in our tests.

APPENDIX C – HOW WE TESTED

Configuring the storage

Creating the volumes

The EMC XtremIO Storage Management Application allows the user to configure and monitor XtremIO storage.

1. Log in to the EMC XtremIO Storage Management Application.
2. Click the Configuration tab.
3. Click Create Volume.
4. Give the volume a name and a size, and click Finish. We used the default Logical Block Size (512) for each volume.
5. Repeat the above steps to create all volumes: OS1, DATA1, DATA2, DATA3, DATA4, and LOG1.

Connecting the hosts to the volumes

Create an initiator group and add initiators to connect the hosts to the storage.

1. Log in to the EMC XtremIO Storage Management Application.
2. Click the Configuration tab.
3. In the left pane, select Initiator Groups.
4. Click Create Initiator Group.
5. Enter a name for the Initiator Group.
6. Click Add.
7. Give the initiator a name.
8. Select ESX as the operating system.
9. Select or enter the initiator port address for the first port of the first host.
10. Repeat for all remaining initiators.
11. Click OK.
12. Click Finish.

Configuring the hosts

Installing the EMC VSI plugin

The EMC Virtual Storage Integrator (VSI) plugin runs as a virtual appliance and handles configuration settings for vSphere hosts attached to EMC storage.

1. Download the vApp from download.emc.com/downloads/DL52899_VSI_for_VMWare_vSphere_Web_Client_6.1.ova
2. Deploy the vApp and set an IP address.
3. Use a web browser to navigate to https://<vApp_IP_address>:8443/vsi_usm/admin and log in with the default credentials of admin/ChangeMe.
4. Change the password if desired.
5. Click Register VSI Plug-in.
6. Enter the vCenter IP address and credentials, and click Register.
7. Navigate to the vCenter Web Client.
8. In the left pane, click Solutions Integration Service.
9. Enter the vApp details, and click OK.
10. Enter the storage array details, and click OK.

Applying the EMC VSI host settings

1. Right-click on a host in the vCenter Web Client.
2. Click EMC VSI Plug-in→Apply Host Settings.
3. Check all checkboxes, and click OK.
4. Reboot the host.
5. Ensure that round-robin pathing is enabled on all storage paths.
6. Apply the EMC VSI settings to the remaining hosts.

Creating the workload virtual machines – DVD Store 2

Creating the VM

1. In VMware vCenter, navigate to Virtual Machines.
2. Click the icon to create a new VM.
3. Leave Create a new virtual machine selected, and click Next.
4. Enter a name for the virtual machine, and click Next.
5. Place the VM on a host with available CPUs, and click Next.
6. Select the OS datastore for the 50GB OS VMDK, and click Next.
7. Click Next.
8. Select the guest OS as Windows Server 2012 R2, and click Next.
9. In the Customize Hardware section, make the following changes:
    a. Increase the vCPUs to 32.
    b. Increase the memory to 4GB.
10. Add a 300GB VMDK for SQL data, and select the VMware Paravirtual controller and Thick Provisioned Eager Zeroed. Place the VMDK in one of the DATA datastores.
11. Add a 250GB VMDK for SQL logs, and select the VMware Paravirtual controller and Thick Provisioned Eager Zeroed. Place the VMDK in the LOGS datastore.
12. Connect the VM to the test network.
13. Click Next.
14. Click Finish.
15. Follow the instructions below to install the guest OS.

Installing Microsoft Windows Server 2012 R2

1. Choose the language, time and currency, and keyboard input. Click Next.
2. Click Install Now.
3. Enter the license key, and click Next.
4. Select Windows Server 2012 R2 Datacenter, and click Next.
5. Accept the license terms, and click Next.
6. Click Custom.
7. Select the OS disk, and click Next.
8. After the installation completes, enter the administrator password twice, and click Finish.

Configuring the operating system

1. Install all available Windows® updates. Restart as necessary.
2. Enable remote desktop access.

3. Bring the SQL data and log VMDKs online, and create NTFS volumes on each.
4. Change the hostname, and reboot when the installer prompts you.
5. Set up networking for the test network:
   a. Click Start→Control Panel, right-click Network Connections, and select Open.
   b. Right-click the management traffic NIC, and select Properties.
   c. Select TCP/IP (v4), and select Properties.
   d. Set the IP address and subnet for the NIC, which will handle client traffic, and click OK.
   e. Click OK, and click Close.

Installing VMware Tools

1. Attach the VMware Tools installation ISO to the VM.
2. Open setup.exe on the installation media, and complete the wizard to install VMware Tools. We used all default settings.

Installing Microsoft SQL Server 2014

1. Attach the installation media to the VM.
2. In the VM guest operating system, run setup.exe.
3. If the installer prompts you with a .NET installation prompt, to enable the .NET Framework Core role, click Yes.
4. In the left pane, click Installation.
5. Click New SQL Server stand-alone installation or add features to an existing installation.
6. At the Setup Support Rules screen, wait for the rule check to complete. If there are no failures or relevant warnings, click OK.
7. Enter your license key, and click Next.
8. Choose to accept the license terms, and click Next.
9. If no failures are displayed after the setup support files are installed, click Next.
10. At the Setup Role screen, choose SQL Server Feature Installation.
11. At the Feature Selection screen, select Database Engine Services, Full-Text and Semantic Extractions for Search, Client Tools Connectivity, Client Tools Backwards Compatibility, Management Tools – Basic, and Management Tools – Complete. Click Next.
12. At the Installation Rules screen, after the check completes, click Next.
13. At the Instance configuration screen, leave the default selection of default instance, and click Next.
14. At the Disk Space Requirements screen, click Next.
15. At the Server Configuration screen, check that NT AUTHORITY\SYSTEM is selected for SQL Server Agent and SQL Server Database Engine. Click Next.
16. Assuming there are no failures, click Next.
17. At the Database Engine Configuration screen, select Mixed Mode.
18. Enter and confirm a password for the system administrator account.
19. Click Add Current user. This may take several seconds.
20. Click the Data Directories tab, and type `E:\` for the default data locations and `F:\` for the default log locations.
21. Click Next.
22. At the Error and usage reporting screen, click Next.

23. At the Installation Configuration Rules screen, check that there are no failures or relevant warnings, and click Next.
24. At the Ready to Install screen, click Install.
25. After installation completes, click Close.

Configuring the database workload client

For our testing, we ran the database workload driver on the SQL VM. To stop and start the workload, and to collect counter results, we created a controller VM running Microsoft Windows Server 2012 R2 and added it to our infrastructure storage array.

Configuring the database

Data generation overview

We generated the data using the Install.pl script included with DVD Store version 2.1 (DS2), providing the parameters for our 250GB database size and the database platform on which we ran: Microsoft SQL Server. We ran the Install.pl script on a utility system running Linux. The Install.pl script also generated the database schema.

After processing the data generation, we transferred the data files and schema creation files to a Windows-based system running SQL Server 2014. We built the 250GB database in SQL Server 2014, and then performed a full backup, storing the backup file on the C: drive for quick access. We used that backup file to restore the server between test runs.

The only modification we made to the schema creation scripts were the specified file sizes for our database. We explicitly set the file sizes higher than necessary to ensure that no file-growth activity would affect the outputs of the test. Besides this file size modification, the database schema was created and loaded according to the DVD Store documentation. Specifically, we followed the steps below:

1. We generated the data and created the database and file structure using database creation scripts in the DS2 download. We made size modifications specific to our 250GB database and the appropriate changes to drive letters.
2. We transferred the files from our Linux data generation system to a Windows system running SQL Server 2014.
3. We created database tables, stored procedures, and objects using the provided DVD Store scripts.
4. We set the database recovery model to bulk-logged to prevent excess logging.
5. We loaded the data we generated into the database. For data loading, we used the import wizard in SQL Server Management Studio. Where necessary, we retained options from the original scripts, such as Enable Identity Insert.
6. We created indices, full-text catalogs, primary keys, and foreign keys using the database-creation scripts.
7. We updated statistics on each table according to database-creation scripts, which sample 18 percent of the table data.
8. On the SQL Server instance, we created a ds2user SQL Server login using the following Transact SQL (TSQL) script:

```
USE [master]
GO
CREATE LOGIN [ds2user] WITH PASSWORD=N'',
         DEFAULT_DATABASE=[master],
```

```
                DEFAULT_LANGUAGE=[us_english],
                CHECK_EXPIRATION=OFF,
                CHECK_POLICY=OFF
GO
```

9. We set the database recovery model back to full.
10. We created the necessary full text index using SQL Server Management Studio.
11. We created a database user and mapped this user to the SQL Server login.
12. We then performed a full backup of the database. This backup allowed us to restore the databases to a pristine state relatively quickly between tests.

Figure 12 shows our initial file size modifications.

Logical name	Filegroup	Initial size (MB)
Database files		
primary	PRIMARY	4
Customer files (8)	DS_CUST_FG	2,560 each
Index files (8)	DS_IND_FG	1,536 each
ds_misc files (8)	DS_MISC_FG	1,536 each
Orders files (8)	DS_ORDERS	1,536 each
Log files		
ds_log	Not Applicable	10,240

Figure 12: Our initial file size modifications.

About running the data transfers

We used copies of the 250GB DS2 database files to serve as our test files for the data transfer. We ran a batch file from the file server to pull the database file from each VM under test.

About running the DVD Store tests

We created a series of batch files, SQL scripts, and shell scripts to automate the complete test cycle. DVD Store outputs an orders-per-minute (OPM) metric, which is a running average calculated through the test. In this report, we report the last OPM reported by each client/target pair.

Each complete test cycle consisted of the following general steps.

1. Clean up prior outputs from the target system and the client driver system.
2. Drop the database from the target.
3. Restore the database on the target.
4. Shut down the target.
5. Reboot the host and client system.
6. Wait for a ping response from the server under test (the hypervisor system), the client system, and the target.
7. Let the test server idle for 20 minutes.
8. Start the DVD Store driver on the client.

We used the following DVD Store parameters for testing the production-level VMs:

```
ds2sqlserverdriver.exe --target=<target_IP> --ramp_rate=10 --
run_time=60 --n_threads=40 --db_size=250GB --think_time=0.01 --
detailed_view=Y --warmup_time=10 --report_rate=1 --csv_output=<drive
path>
```

We used the following DVD Store parameters for testing the development-level VMs:

```
ds2sqlserverdriver.exe --target=<target_IP> --ramp_rate=10 --
run_time=60 --n_threads=22 --db_size=250GB --think_time=0.01 --
detailed_view=Y --warmup_time=10 --report_rate=1 --csv_output=<drive
path>
```

Creating the workload virtual machines – SLOB

Creating the VM

1. In VMware vCenter, navigate to Virtual Machines.
2. Click the icon to create a new VM.
3. Leave Create a new virtual machine selected, and click Next.
4. Enter a name for the virtual machine, and click Next.
5. Place the VM on a host with available CPUs, and click Next.
6. Select the OS datastore for the 50GB OS VMDK, and click Next.
7. Click Next.
8. Select the guest OS as Oracle Enterprise Linux 6, and click Next.
9. In the Customize Hardware section, make the following changes:
    a. Increase the vCPUs to 8.
    b. Increase the memory to 4GB.
10. Add a 50GB VMDK for Oracle data, and select the VMware Paravirtual controller and Thick Provisioned Eager Zeroed. Place the VMDK in one of the DATA datastores.
11. Repeat the previous step to add a 50GB VMDK to each DATA datastore, for four data VMDKs.
12. Add a 50GB VMDK for Oracle logs, and select the VMware Paravirtual controller and Thick Provisioned Eager Zeroed. Place the VMDK in the LOGS datastore.
13. Connect the VM to the test network.
14. Click Next.
15. Click Finish.
16. Follow the instructions below to install the guest OS.

Installing Oracle Enterprise Linux 6.5

1. Attach the installation ISO to the VM, and boot to it.
2. Select Install or upgrade an existing system.
3. If you are unsure of the fidelity of the installation disk, select OK to test the installation media; otherwise, select Skip.
4. In the opening splash screen, select Next.
5. Choose the language you wish to use, and click Next.
6. Select the keyboard layout, and click Next.
7. Select Basic Storage Devices, and click Next.

8. Select Fresh Installation, and click Next.
9. Insert the hostname, and select Configure Network.
10. In the Network Connections menu, configure network connections.
11. After configuring the network connections, click Close.
12. Click Next.
13. Select the nearest city in your time zone, and click Next.
14. Enter the root password, and click Next.
15. Select Use All Space, and click Next.
16. When the installation prompts you to confirm that you are writing changes to the disk, select Write changes to disk.
17. Select Software Basic Server, and click Next. Oracle Enterprise Linux installation begins.
18. When the installation completes, select Reboot to restart the server.

Installing VMware Tools

1. Right-click the VM in the Web Client, and select Install/Upgrade VMware Tools.
2. Log on to the guest as root.
3. Mount the CD ROM device:
   ```
   mount /dev/cdrom /mnt
   ```
4. Untar VMware Tools into a temporary directory:
   ```
   tar –C /tmp –zxf /mnt/VMwareTools-9.4.0-1280544.tar.gz
   ```
5. Run the install script and accept the defaults:
   ```
   /tmp/vmware-tools-distrib/vmware-install.pl
   ```
6. Follow the prompts to configure and install VMware tools. The installer will automatically load the NIC drivers, create a new initrd, and unmount the CD.
7. Reboot the VM.

Initial configuration tasks

Complete the following steps to provide the functionality that Oracle Database requires. We performed all of these tasks as root.

1. Disable firewall services. In the command line (as root), type:
   ```
   service iptables stop
   chkconfig iptables off
   service ip6tables stop
   chkconfig ip6tables off
   ```
2. Edit `/etc/selinux/config`:
   ```
   SELINUX=permissive
   ```
3. Modify `/etc/hosts` to include the IP address of the internal IP and the hostname.
4. Edit `/etc/security/limits.d/90-nproc.conf`:
   Change this:
   ```
   soft nproc 1024
   ```
   To this:
   ```
   - nproc 16384
   ```
5. Install 12*c* RPM packages, resolve package dependencies, and modify kernel parameters:
   ```
   yum install oracle-rdbms-server-12cR1-preinstall –y
   ```

6. Install automatic system tuning for database storage through yum:

```
yum install tuned
chkconfig tuned on
tuned-adm profile enterprise-storage
```

7. Using yum, install the following prerequisite packages for Oracle Database:

```
yum install elfutils-libelf-devel
yum install xhost
yum install unixODBC
yum install unixODBC-devel
yum install oracleasm-support oracleasmlib oracleasm
```

8. Create the `oracle` user account and groups and password:

```
groupadd -g 1003 oper
groupadd -g 1004 asmadmin
groupadd -g 1005 asmdba
groupadd -g 1006 asmoper
usermod -G dba,oper,asmadmin,asmdba,asmoper oracle
passwd oracle
```

9. Create the `/u01` directory for Oracle inventory and software and give it to the Oracle user:

```
mkdir -p /u01/app/oracle/product/12.1.0/grid_1
mkdir -p /u01/app/oracle/product/12.1.0/dbhome_1
chown -R oracle:oinstall /u01
chmod -R 775 /u01
```

10. Edit bash profiles to set up user environments:

```
# vim /home/oracle/.bash_profile
# Oracle Settings
   export TMP=/tmp
   export TMPDIR=$TMP
   export ORACLE_HOSTNAME=localhost.localdomain
   export ORACLE_BASE=/u01/app/oracle
   export GRID_HOME=$ORACLE_BASE/product/12.1.0/grid_1
   export DB_HOME=$ORACLE_BASE/product/12.1.0/dbhome_1
   export ORACLE_HOME=$DB_HOME
   export ORACLE_SID=orcl
   export ORACLE_TERM=xterm
   export BASE_PATH=/usr/sbin:$PATH
   export PATH=$ORACLE_HOME/bin:$BASE_PATH
   export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
   export
   CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib
   alias grid_env='. /home/oracle/grid_env'
   alias db_env='. /home/oracle/db_env'

# vim /home/oracle/grid_env
```

```
export ORACLE_SID=+ASM1
export ORACLE_HOME=$GRID_HOME
export PATH=$ORACLE_HOME/bin:$BASE_PATH
export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
export
CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib
```

```
# vim /home/oracle/db_env

export ORACLE_SID=orcl
export ORACLE_HOME=$DB_HOME
export PATH=$ORACLE_HOME/bin:$BASE_PATH
export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
export
CLASSPATH=$ORACLE_HOME/JRE:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib
```

Adding the log and four data drives to the VM and modifying the vmfx file

1. Power off the VM.
2. Right-click the VM in the vSphere Web Client, and choose Edit Settings…
3. Click on the VM Options tab, and expand the Advanced menu option.
4. Choose Edit Configuration…
5. Click Add Row, and type `disk.EnableUUID` in the parameter field and `TRUE` in the value field.
6. Go back to the Virtual Hardware tab.
7. Click the drop-down menu for New device, and choose New Hard Disk.
8. Name the Hard Disk and choose the size that you want it to be.
9. Repeat steps 7 and 8 for all remaining drives.
10. Click OK.
11. Power the VM back on.

Configuring disks for ASM

1. For each of the five shared disks, create a GPT label, and create one partition. For example, see the following shell script:
```
for disk in sdc sdd sde sdf sdg; do
parted /dev/$disk mklabel gpt
parted /dev/$disk mkpart primary "1 -1"
done
```
2. If desired, label the disk's partition with its Oracle function. For example:
```
parted /dev/sdc name 1 DATA1
parted /dev/sdd name 1 DATA2
parted /dev/sde name 1 DATA3
parted /dev/sdf name 1 DATA4
parted /dev/sdg name 1 LOG1
```
3. Initialize Oracle ASM on each server by executing the following commands as root on each node:
```
oracleasm init
oracleasm configure -e -u grid -g oinstall -s y -x sda
```
4. Label each shared disk-partition with an appropriate ASM name. For example, following the OS partition names created above, execute the following commands on one system:

```
oracleasm createdisk DATA1 /dev/sdc1
oracleasm createdisk DATA2/dev/sdd1
oracleasm createdisk DATA3 /dev/sde1
oracleasm createdisk DATA4 /dev/sdf1
oracleasm createdisk LOG1 /dev/sdg1
```

5. On each server, scan the disks to make the disks immediately available to Oracle ASM.

```
oracleasm scandisks
oracleasm listdisks
```

Installing Oracle Grid Infrastructure 12*c*

1. Log in as the `oracle` user.
2. Unzip linuxamd64_12c_grid_1of2.zip and linuxamd64_12c_grid_2of2.zip
3. Open a terminal to the unzipped database directory.
4. Type `grid_env` to set the Oracle grid environment.
5. To start the installer, type `./runInstaller`
6. At the Updates screen, select Skip updates.
7. In the Select Installation Option screen, select Install and Configure Grid Infrastructure for a Standalone Server, and click Next.
8. Choose the language, and click Next.
9. In the Create ASM Disk Group screen, choose the Disk Group Name, and change redundancy to External.
10. Select the four disks that you are planning to use for the database, and click Next.
11. In the Specify ASM Password screen, choose Use same password for these accounts, write the passwords for the ASM users, and click Next.
12. Leave the default Operating System Groups, and click Next.
13. Leave the default installation, and click Next.
14. Leave the default inventory location, and click Next.
15. Under Root script execution, select Automatically run configuration scripts and enter root credentials.
16. In the Prerequisite Checks screen, make sure that there are no errors.
17. In the Summary screen, verify that everything is correct. To install Oracle Grid Infrastructure click Finish.
18. At one point during the installation, the installation prompts you to execute two configuration scripts as root. Follow the instructions to run the scripts.
19. At the Finish screen, click Close.
20. To run the ASM Configuration Assistant, type `asmca`.
21. In the ASM Configuration Assistant, click Create.
22. In the Create Disk Group window, name the new disk group `log`, choose redundancy External (None), and select the log disk for redo logs.
23. Click Advanced Options, and type `12.1.0.0.0` in ASM Compatibility and Database Compatibility. Click OK.
24. Right-click the DATA drive, and choose Edit Attributes. Make sure both ASM and Database Compatibility fields list `12.1.0.0.0`, and click OK.
25. Exit the ASM Configuration Assistant.

Installing Oracle Database 12*c*

1. Unzip linuxamd64_12c_database_1_of_2.zip and linuxamd64_12c_database_2_of_2.zip.
2. Open a terminal to the unzipped database directory.
3. Type `db_env` to set the Oracle database environment.

4. Run `./runInstaller.sh`.
5. Wait until the GUI installer loads.
6. On the Configure Security Updates screen, enter the credentials for My Oracle Support. If you do not have an account, uncheck the box I wish to receive security updates via My Oracle Support, and click Next.
7. At the warning, click Yes.
8. On the Download Software Updates screen, enter the desired update option, and click Next.
9. On the Select Installation Option screen, select Install database software only, and click Next.
10. On the Grid Installation Options screen, select Single instance database installation, and click Next.
11. On the Select Product Languages screen, leave the default setting of English, and click Next.
12. On the Select Database Edition screen, select Enterprise Edition, and click Next.
13. On the Specify Installation Location, leave the defaults, and click Next.
14. On the Create Inventory screen, leave the default settings, and click Next.
15. On the Privileged Operating System groups screen, keep the defaults, and click Next.
16. Allow the prerequisite checker to complete.
17. On the Summary screen, click Install.
18. Once the Execute Configuration scripts prompt appears, ssh into the server as `root`, and run the following command:

```
/home/oracle/app/oracle/product/12.1.0/dbhome_1/root.sh
```

19. Return to the prompt, and click OK.
20. Once the installer completes, click Close.

Adding hugepages

In Putty, edit `/etc/sysctl.conf` by adding the following lines:

```
vm.nr_hugepages = 13312
vm.hugetlb_shm_group = 1001
```

Creating and configuring the database

1. Using Putty with X11 forwarding enabled, SSH to the VM.
2. Type `dbca`, and press Enter to open the Database configuration assistant.
3. At the Database Operation screen, select Create Database, and click Next.
4. Under Creation Mode, select Advanced Mode, and click Next.
5. At the Select Template screen, select General Purpose or Transaction Processing. Click Next.
6. Enter a Global database name and the appropriate SID.
7. At the Management Options screen, select Configure Enterprise Manager (EM) Database Express. Click Next.
8. At the Database Credentials screen, select Use the Same Administrative Password for All Accounts. Enter a password, and click Next.
9. At the Network Configuration screen, click Next.
10. At the Storage Configuration screen, select Automatic Storage Management, and select `+DATA` as the Database.
11. At the Database Options screen, click Next.
12. At the Initialization Parameters screen, click use Automatic Memory Management.
13. At the Creation Options screen, select Create Database, and click Next.
14. At the summary screen, click Finish.
15. Close the Database Configuration Assistant.

16. In a web browser, browse to `https://vm.ip.address:5500/em` to open the database manager.
17. Log in as system with the password you specified.
18. Go to Storage→Tablespaces.
19. Click Create.
20. Enter `SLOB` as the Name, and check the Set As Default box. Click OK.
21. Go to Storage→Redo Log Groups.
22. Click Actions→Switch file… until you get one of the groups to go inactive.
23. Highlight the inactive group, and click Actions→Drop group.
24. In Putty, enter sqlplus and type ALTER DATABASE ADD LOGFILE GROUP 3 ('+logs/orcl/onlinelog/redox.log') SIZE 9G BlOCKSIZE 4k; where x is the number of the log file (1, 2, etc.).
25. Repeat steps 21 – 24 until all default logs are dropped, and two new ones exist.
26. Click on Configuration→Initialization Parameters
27. Change the `use_large_pages` parameter to Only.

Installing SLOB and populating the database

1. Download the SLOB kit from www.kevinclosson.net/slob/.
2. Copy and untar the files to `/home/oracle/dev/`.
3. Type `./setup.sh SLOB 128` to start the data population to the SLOB tablespace we created earlier.
4. When the setup is complete, the database is populated.

Running the performance test

To run the SLOB workload, we first rebooted all hosts and VMs and waited 20 minutes to allow the servers to reach full idle. We ran the production-level workload with `./runit.sh 70`, 0.1ms think time, and the development-level workloads with `./runit.sh 128`, 0.7ms think time.

Creating the workload virtual machines – DB2 IOPS

Creating the VM

1. In VMware vCenter, navigate to Virtual Machines.
2. Click the icon to create a new VM.
3. Leave Create a new virtual machine selected, and click Next.
4. Enter a name for the virtual machine, and click Next.
5. Place the VM on a host with available CPUs, and click Next.
6. Select the OS datastore for the 50GB OS VMDK, and click Next.
7. Click Next.
8. Select the guest OS as Red Hat Enterprise Linux 6, and click Next.
9. In the Customize Hardware section, make the following changes:
    a. Increase the vCPUs to 8.
    b. Increase the memory to 4GB.
10. Add a 100GB VMDK for DB2 data, and select the VMware Paravirtual controller and Thick Provisioned Eager Zeroed. Place the VMDK in one of the DATA datastores.
11. Repeat the previous step to add a 100GB VMDK to another DATA datastore, for two data VMDKs.
12. Add a 100GB VMDK for DB2 logs, and select the VMware Paravirtual controller and Thick Provisioned Eager Zeroed. Place the VMDK in the LOGS datastore.

13. Connect the VM to the test network.
14. Click Next.
15. Click Finish.
16. Follow the instructions below to install the guest OS.

Installing Red Hat Enterprise Linux 6.6

1. Attach the installation media ISO to the VM, and boot to it.
2. Select Install or upgrade an existing system.
3. If you are unsure of the fidelity of the installation disk, to test the installation media, select OK; otherwise, select Skip.
4. In the opening splash screen, select Next.
5. Choose the language you wish to use, and click Next.
6. Select the keyboard layout, and click Next.
7. Select Basic Storage Devices, and click Next.
8. Select Fresh Installation, and click Next.
9. Insert the hostname, and select Configure Network.
10. In the Network Connections menu, configure the adapter to connect to the testing network.
11. After configuring the network connections, click Close.
12. Click Next.
13. Select the nearest city in your time zone, and click Next.
14. Enter the root password, and click Next.
15. Select Replace Existing Linux file system, and check the Review the file system checkbox.
16. Edit the file system as follows:
```
lv_root 30216 (ext4)
lv_swap 20480 (swap)
```
17. Format and write changes to the disk.
18. Select the Software Development Workstation package.
19. When the installation completes, to restart the server, select Reboot.
20. Upon reboot:
    a. Agree to the licensing terms.
    b. Choose to register later.
    c. Create a new user named `user`
    d. Configure NTP.
21. Leave default settings for Kdump.
22. Reboot again.

Installing VMware Tools

1. Right-click the VM in the Web Client, and select Install/Upgrade VMware Tools.
2. Log on to the guest as `root`.
3. Mount the CD ROM device:
```
mount /dev/cdrom /mnt
```
4. Untar VMware Tools into a temporary directory:
```
tar -C /tmp -zxf /mnt/VMwareTools-9.4.0-1280544.tar.gz
```
5. Run the install script and accept the defaults:

```
/tmp/vmware-tools-distrib/vmware-install.pl
```

6. Follow the prompts to configure and install VMware tools.
7. The installer will automatically load the NIC drivers, create a new initrd, and unmount the CD.
8. Reboot the VM.

Initial configuration tasks

Complete the following steps to provide the functionality that DB2 requires. We performed all of these tasks as root.

1. Disable firewall services. In the command line (as root), type:

```
service iptables stop
chkconfig iptables off
service ip6tables stop
chkconfig ip6tables off
```

2. Edit /etc/selinux/config:

```
SELINUX=permissive
```

3. Modify `/etc/hosts` to include the IP address of the internal IP and the hostname.

Storage configuration

Make the following adjustments to the data and log disks.

1. Log in to the DB2 VM as the root user, and run the following commands:

```
parted /dev/sdb mklabel gpt
parted /dev/sdb mkpart primary "1 -1"
parted /dev/sdb name 1 data1
mkfs.ext4 /dev/sdb1
parted /dev/sdc mklabel gpt
parted /dev/sdc mkpart primary "1 -1"
parted /dev/sdc name 1 data2
mkfs.ext4 /dev/sdc1
parted /dev/sdd mklabel gpt
parted /dev/sdd mkpart primary "1 -1"
parted /dev/sdd name 1 log
mkfs.ext4 /dev/sdd1
mkdir /mnt/data1
mkdir /mnt/data2
mkdir /mnt/log
mkdir /mnt/cifs
chmod 777 /mnt/cifs
```

2. Edit `/etc/fstab` to include the following lines:

```
/dev/sdb1 /mnt/data1  ext4  defaults   1 1
/dev/sdc1 /mnt/data2  ext4  defaults   1 1
/dev/sdd1 /mnt/log    ext4  defaults   1 1
//<gateway_server_ip_address>/install /mnt/cifs cifs
_netdev,username=administrator,password=Password1,dir_mode=0755,file_m
ode=0755,uid=500,gid=500 0 0
```

Installing DB2

1. Make the directories /home/db2software and /home/db2iops.
2. Copy the installation software into the /home/db2software directory.
3. Extract the installation files.

4. Run /home/db2software/server/db2install to begin the installation process.
5. In the left pane, select Install a product, and select Install new.
6. Select DB2 Version 10.5, and click Install new.
7. At the welcome screen for DB2, click Next.
8. Accept the license agreement, and click Next.
9. Select Typical.
10. Click Install DB2 Server Edition on this computer.
11. Enter /opt/ibm/db2/V10.5 as the installation directory, and click Next.
12. Enter a password for dasusr1, and click Next.
13. Click Create DB2 instance, and click Next.
14. Click Single partition instance, and click Next.
15. Enter a password for db2inst1, and click Next.
16. Enter a password for db2fenc1, and click Next.
17. Disable setup notifications, and click Next.
18. Click Finish.

Installing FP5

1. From the /home/db2software directory, extract v10.5fp5_linuxx64_universal_fixpack.tar, and then run the following commands:
```
su db2inst1
db2stop
exit
su dasusr1
db2admin stop
exit
```
2. Run ./universal/installFixPack to install FP4.
3. Enter the DB2 path as /opt/ibm/db2/V10.5
4. Enter no to keep the default installation directory.
5. Copy the installation software into the /home/db2software directory.

Licensing the server

Run the following commands to apply the DB2 license:
```
su db2inst1
db2start
db2 licm -a /home/db2software/License_Key/db2/Licensedb2aese_c.lic
```

Installing DB2IOPS

Run the following to commands to install the DB2IOPS benchmark:
```
mkdir /home/db2iops
cp /<source directory>/db2iops.tar /home/db2iops
cd /home/db2iops
tar -xvf db2iops.tar
./install.sh
```

Configuring DB2IOPS

Run the following commands to enable full permissions on the data and log disks:
```
chmod 777 /mnt/data1/*
chmod 777 /mnt/data2/*
chmod 777 /mnt/log/*
```

```
chmod 777 /home/db2ins1/*
```

Run `visudo` and add the following line to the end:

```
db2inst1 ALL=(ALL) NOPASSWD:ALL
```

Open `/home/db2iops/db2_iops.cfg` and edit the buffer pool size. We used a buffer pool size of 250 for the production-level VMs and a buffer pool size of 350 for the development-level VMs.

Running DB2IOPS

To run DB2IOPS, open `/home/db2iops/db2_iops.sh` and press `1` to create a database. Once the database is created, press `2` to run a workload against it. Once the run is complete, press `4` to delete the database and press `1` to create a fresh copy of the database.

APPENDIX D – DETAILED RESULTS

Figure 13 shows the total IOPS during our median test run.

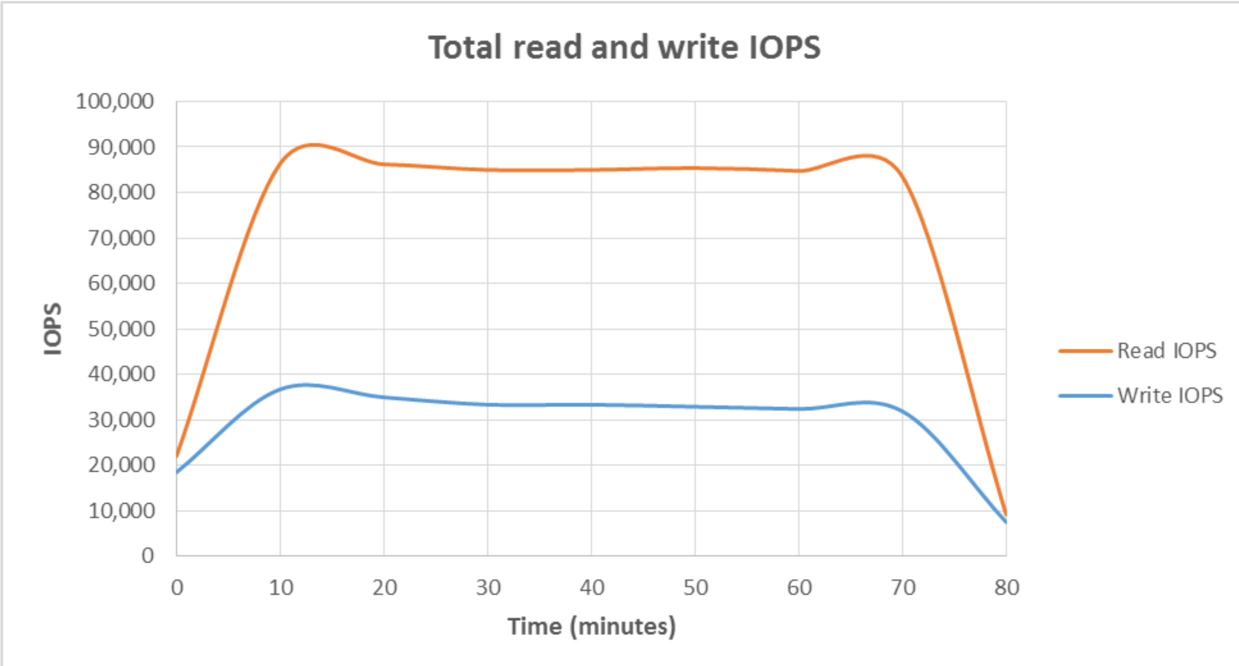

Figure 13: The total number of read and write IOPS during our median 80-minute test.

Figure 14 shows the average latency during our median test run.

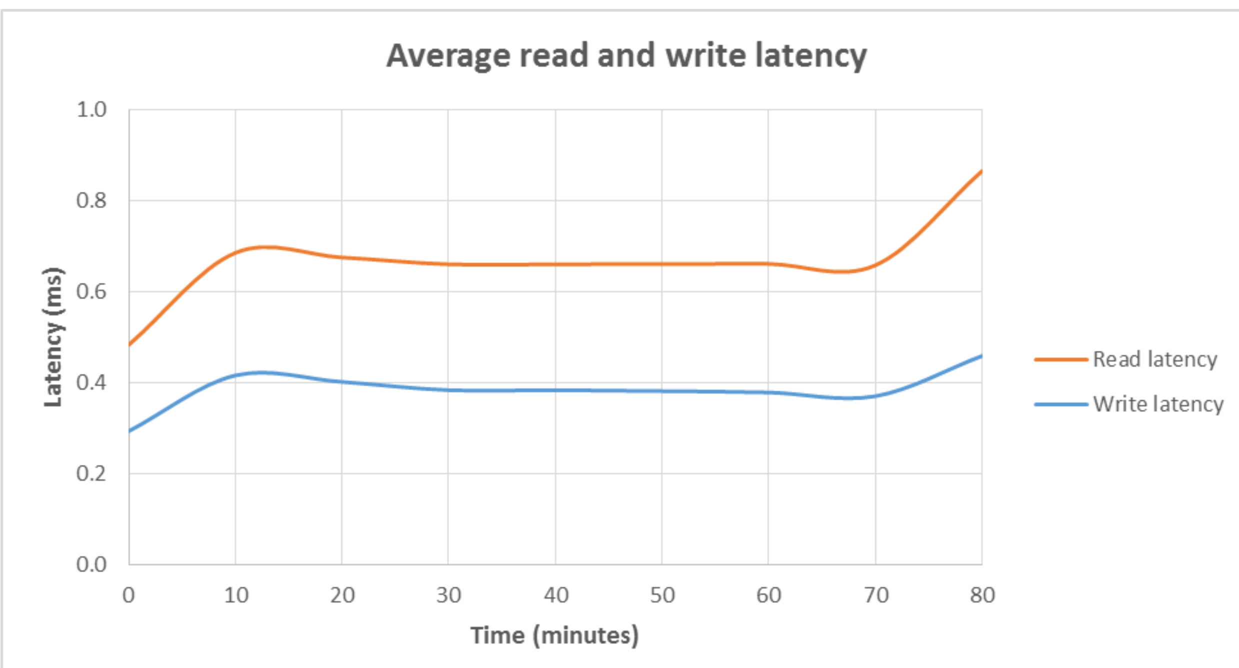

Figure 14: The average latency of read and write IOPS during our median 80-minute test.

ABOUT PRINCIPLED TECHNOLOGIES

Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.