

Evaluation Report: HPE MSA 2040 Storage in a Mixed Workload Virtualized Environment

Evaluation report prepared under contract with Hewlett Packard Enterprise (HPE)

Executive Summary

Enterprise computing hardware can be expensive. Mitigating this cost has been a large part of the trend to increasingly virtualized computing environments. Hardware vendors recognize this and understand that customers often want to have the ability to run multiple applications on a single piece of hardware, such as a storage array, rather than buying multiple devices, each dedicated to a single purpose. Hard disk drive (HDD) only arrays can have trouble supporting diverse workloads with differing I/O patterns without requiring a lot of drive spindles. This often results in short-stroking (data written only to the outer sectors of the disk's platters), significantly over-provisioning the number of drives, wasting a good deal of the drive space to keep IOPS and bandwidth up, whereas a single solid state drive can deliver the IOPS and bandwidth of an entire shelf or more of HDDs.

HPE's MSA 2040 Storage array is a flexible hybrid array ideal for small and medium businesses seeking to cut costs through application consolidation. Understanding that these businesses demand the same return on investment as large enterprises, HPE offers upgrade options to the MSA 2040 Storage by allowing customers to convert an all HDD array into a hybrid array by replacing some of the HDDs with solid state drives (SSD) to accelerate performance without adding expansion shelves. With the addition of a customer-defined number of SSDs, the array can be configured to accelerate read performance with an SSD read cache, or read and write performance with HPE's Performance Tier™ performance tiering solution.

HPE commissioned Demartek to evaluate the suitability of the MSA 2040 Storage to support a mixed application workload in a VMware vSphere environment. Running common applications like Microsoft Exchange Server, SQL Server, SharePoint, a web server and an SMB fileserver simultaneously we compared the performance of an array with all HDD storage to that with an SSD read cache and an SSD Performance Tier configuration.

For the workload mixture we deployed, we saw storage performance improvement with both acceleration technologies, to a maximum 2.4X bandwidth improvement with SSD Performance Tier configured. Storage latencies decreased by up to 70%, resulting in a dramatic 5X reduction in the database application storage latency.

The HPE MSA 2040 Storage

The HPE MSA 2040 Storage Array (Figure 1) is a dual controller, two rack unit, 24 small form factor drive storage array as part of HPE's MSA product line. The MSA 2040 model supports 8 Gb and 16 Gb Fibre Channel along with 1 Gb and 10 Gb iSCSI. Each array can support seven expansion units for up to 768 TB of total storage capacity.



Figure 1 - HPE MSA 2040 Storage

The base configuration of 24 HDDs is upgradable by replacing some HDDs with flash storage. The MSA's Advanced Virtualization features, available in GL200 or later MSA firmware, manages the addition of SSDs to an array by enabling the creation of an SSD read-cache, or configuring it as a full R/W tier using the MSA Performance Tiering SW License. Both features use a real-time sub-LUN tiering engine which migrates active data between HDDs and SSDs every five seconds.

For this evaluation, HPE provided Demartek with an MSA 2040 array with 20 900 GB 10K RPM HDDs and four 400 GB SSDs. With this hardware, we were able to set up a 20 HDD baseline test case, with 10 drives provisioned to each controller in a RAID 6 configuration. Then to evaluate read caching, we added a single SSD to each controller and assigned those SSDs as read cache type in the virtual disk group. For measuring the effect of the optional Performance Tiering license, two SSDs in a RAID 1 disk group were added to each controller.

Three volumes were created from these drives. One volume was used for all virtual machine files and the other two, much larger volumes, were dedicated as datastores for application data. The application virtual drives were provisioned evenly from disk space managed by each controller to best spread the I/O load across the array. Appendix B diagrams the array configuration and volume mapping to vSphere datastores.

The Mixed Workload

Recognizing that cost conscious businesses look to get the best return on their hardware investments, we considered how the MSA 2040 Storage array would support a vSphere environment with several application servers running in virtual machines (VMs). We envisioned a hypothetical small to medium sized business of about 150 employees and created several VMs to support common applications. All VMs shared the same MSA 2040 Storage array for virtual machine files and application data. These included infrastructure VMs for VMware vCenter server and Microsoft Active Directory server. Five application VMs were then added to vSphere. Setup details for the virtual machines are described in Appendix A.

We did not design this workload to try and generate performance bounding benchmarks or specifically highlight any particular I/O patterns. Our intention was to model what seemed a reasonable use case for the business we envisioned. This included some number of the users exercising different applications at the same time. As with real businesses, only a small number of users are likely to be accessing any one IT resource at a given time. We made some educated guesses as to what those user counts could potentially be.

The workload included VMs hosting Microsoft Exchange, SQL Server, and SharePoint, along with a static content Webserver and an SMB fileserver. More detailed descriptions of the workloads and user counts are available in appendix C.

Performance Metrics

Key metrics for storage system performance analysis are I/Os per second (IOPS), bandwidth, and latency or response time. These metrics are defined as follows:

- ◆ **IOPS** – I/Os per second – a measure of the total I/O operations (reads and writes) issued by the application servers.
- ◆ **Bandwidth** – a measure of the data transfer rate, or I/O throughput, measured in MegaBytes per second (MBPS).
- ◆ **Latency** – a measure of the time taken to complete an I/O request, also known as response time. This is frequently measured in milliseconds (one thousandth of a second). Latency is introduced into the SAN at many points, including the server and HBA, SAN switching, and at the storage target(s) and media.

It is important to consider all three metrics when evaluating the performance of storage systems because all three contribute to how the storage will support an application. IOPS

drive bandwidth. The number of IOPS times the I/O request size determines the amount of bandwidth delivered.

Latency can have a very significant effect on application performance and user experience. Transactional applications can be quite sensitive. For instance, Microsoft Exchange Jetstress sets an upper limit of 20 milliseconds for database reads and writes to certify any storage system. Unlike IOPS and bandwidth, where more is better, with latency the goal is to keep it as low as possible.

Performance metrics were gathered at the hypervisor, with the vSphere esxtop utility, and within the VMs with Windows Perfmon and Linux systats.

Results and Analysis

This multi-application workload is not a stress testing benchmark. The goal was to model a virtualized environment and the type of work representative of a small to mid-sized business. All businesses will be different with respect to the exact type and intensity of storage I/O, and therefore these results should be evaluated as an example of one possible scenario.

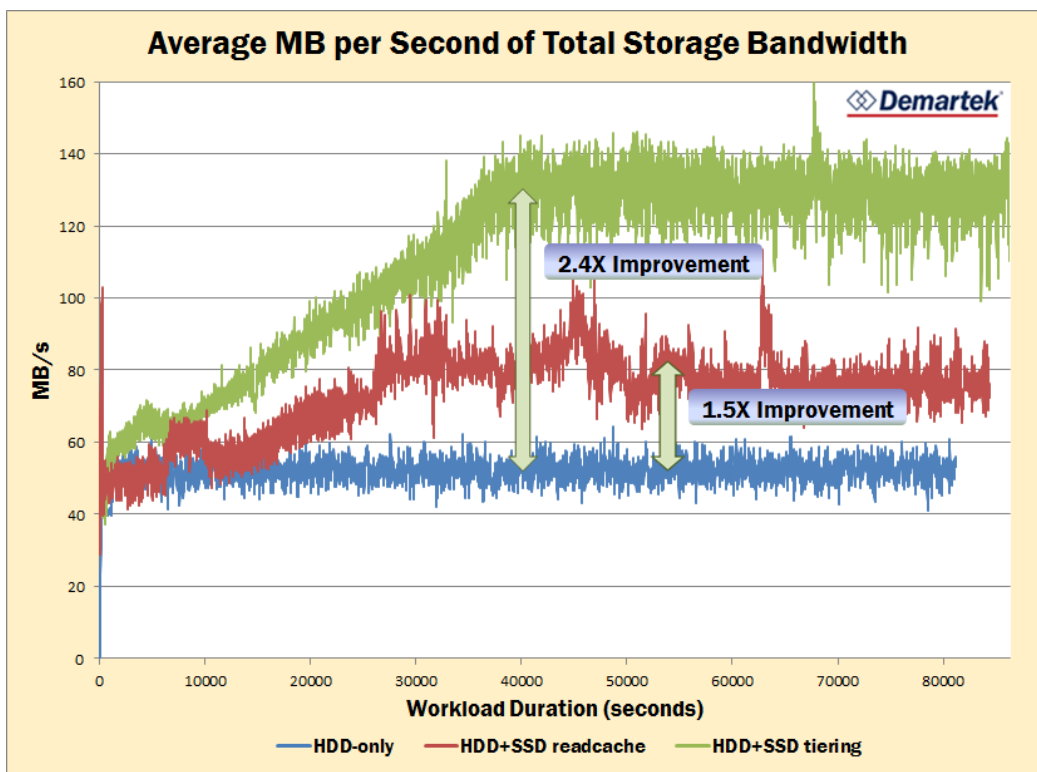


Figure 1 - Total Bandwidth

This is a complex and punishing workload to run on a single array. The mixtures of applications chosen have I/O patterns that cause competition for storage resources, such as HDD head motions, cache space, and controller performance. Going into the testing, we expected the HDD-only array to be challenged to manage all the I/O. Without flash enhancement, the HDD-only array was able to deliver about 52 MB/s of total bandwidth (Figure 1), which included all VM files and vSphere infrastructure, and all application data. At first glance, latency looks pretty good for an all-HDD array, at an average of 7.1 milliseconds when measured at the hypervisor layer (Figure 2).

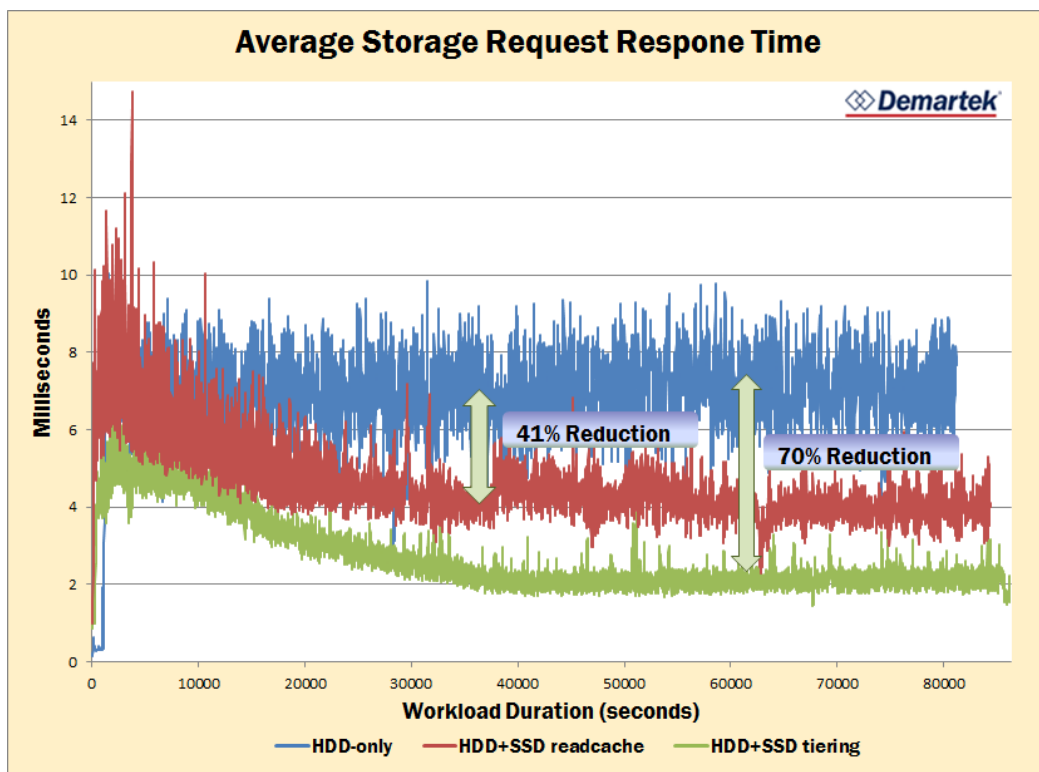


Figure 2 - Storage system response times

However, storage array response times don't tell the whole story. A low I/O latency for operations delivered from the hypervisor to the array doesn't mean that virtual drives provisioned from a vSphere datastore behave the same way. Array performance contributes to the overall picture, but there is overhead in the hypervisor processes, not to mention the guest OS, when managing virtualized storage. Obviously, the better the array latency is, the more benefit can be passed on to the hypervisor and guests. To get a clearer picture of the application experience we took a look at the response times of the virtual drives provisioned to VMs. These response times are always higher than the storage device latency.

Deeming Jetstress the most critical application, since it defines its own storage certification criteria, we determined that if any array configuration couldn't meet the Jetstress requirements we would consider that configuration unsuitable for our workload. We found that the highest average database transactional latency experienced by Microsoft Jetstress was 17.7 milliseconds (Figure 3), which was within the threshold of acceptability set by Jetstress. As such, every configuration, with and without SSD acceleration, met our initial criteria.

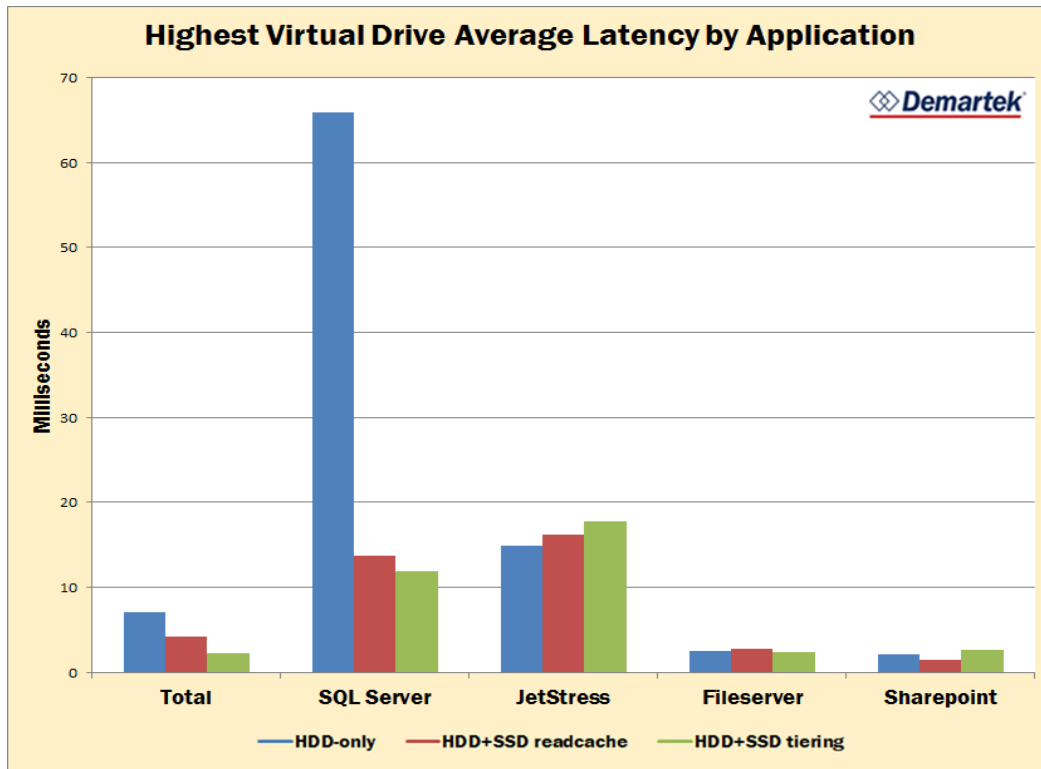


Figure 3 - Virtual drive latencies

SQL Server was the overall biggest winner with regard to improvement when either flash solution was implemented, with latency decreased by roughly a factor of 5 in both cases. The remaining applications did not seem to benefit significantly when comparing latency before and after flash enhancements. This may be due to how the ratio of read-to-write I/O in each application was addressed by the caching and tiering algorithms, or the fact that this workload mix was simply trying to model a realistic use case and these individual workloads just didn't require that much acceleration to deliver the performance needed to run them. This latter case is certainly true for Jetstress as it is defined to deliver a specific performance profile. Curiously, Jetstress, the fileserver, and SharePoint saw latency slightly increase when the array was upgraded with SSDs, in one or both of the

configurations. We suspect this is related to SQL Server's strong performance boost. As the database workload runs faster, its transactions require more data. With limited SSD space some portion of this data must come from the HDDs, which drives up the number of I/O requests. However, the HDDs cannot serve any additional I/O over the baseline, so queuing occurs. As more I/Os are queued on the HDDs, latency goes up. The impact may be very small on the raw devices, but the added overhead of the hypervisor and guest OS magnify it as seen here.

The same appears to be true with regard to application bandwidth. Once we enhanced the MSA 2040 Storage array with SSDs and configured with either read caching or Performance Tier the combined bandwidth of all applications improved. Performance Tier saw the best improvement, 2.4 times the HDD-only configuration, to nearly 130 MB/s while response times dropped by 70% to roughly 2 milliseconds (Figures 1 and 2). Read caching saw more modest gains, but still provided measurable bandwidth and response time improvements.

We noticed that, as defined, the Webserver workload was quickly cached into its virtual machine RAM, performing almost no I/O that reached the drives. Therefore we focused on analyzing the storage bandwidth of the four remaining applications, although SharePoint also appeared to be very RAM cache friendly and did very little drive I/O as well, as the "heavy-hitters" (Figure 4). As with latency, the SQL Server workload again saw the greatest benefit, with total storage bandwidth increasing from 12 MB/s to 39 MB/s to 88MB/s with read caching and Performance Tier respectively. This produced 196 database transactions/second with only HDDs, 409 transactions/second through read caching, and 900 transactions/second with Performance Tier




Figure 4- Application bandwidths

It seems clear that this workload mix favored Performance Tier as a solution and the fairly random I/O profile of the SQL Server OLTP workload took the best advantage of the SSDs. We would speculate that the I/O profiles of the other workloads were both neither intense nor random enough to meet the algorithm's criteria for caching or migrating data to SSD. However, neither were they negatively impacted by the SQL Server workload's performance gains. If we assume they were meeting an SLA prior to the array upgrade, those workloads continued to meet that SLA afterwards.

Summary and Conclusion

Consolidation of workloads into a virtualized environment can drive up the return on investment for computing hardware, but it needs to be approached with a strong understanding of the workload I/O demands. Combining applications that compete for storage resources can be punishing to a drive array. Different businesses will have different requirements that demand different solutions. HPE's MSA 2040 Storage offers three solutions--all HDD, SSD read caching, and SSD performance tiering(aka Performance Tier)-- in one array product.

So long as the business SLA expectations are met, a solution can generally be considered appropriate for the use case. Using our workload as an example, we declared upfront that Exchange Jetstress thresholds had to be met for the storage to be validated. This condition was satisfied in all three configurations. When we enabled flash acceleration, we saw database storage bandwidths improve dramatically, by a factor of 3.5 when read caching, and a factor of more than 7 with Performance Tier. While the database application was experiencing such significant improvement, the remaining application workloads experienced no degradation.

There is also a cost savings to consider. An all HDD solution would have required 48 drives to duplicate the performance of 20 HDDs and four SSDs configured for Performance Tier.¹ That is another entire drive shelf, with all the associated costs for purchase, deployment, and support.

While we would suggest that SSD performance tiering is the better solution for our test case, a business with a different set of applications might find read caching a stronger solution, particularly if their I/O profile included heavier read activity, with potential "hotspots" or frequently accessed data. The great strength of the 2040 Storage array is that it is easily configured for either option and the number of SSDs can be tailored to fit a business' individual needs.

A successful business is not likely to see application usage go down over time and wise IT management will consider anticipated storage growth and performance needs when planning purchases and upgrades. The HPE MSA 2040 Storage provides the benefits of

¹ With four rather small SSDs configured for Performance Tier, aggregate performance increased 2.4 times (bandwidth and IOPS). If the relationship of spindles to performance is linear, another 28 HDDs would be needed to achieve the same performance values as with Performance Tier. (i.e., 20 HDDs x 2.4 = 48 HDDs)



flash storage, to support current needs and to future-proof the enterprise, without the expense of an all-flash array or the inconvenience and waste of short-stroking multiple HDD shelves to achieve the performance demanded.

Appendix A – Test Description and Environment

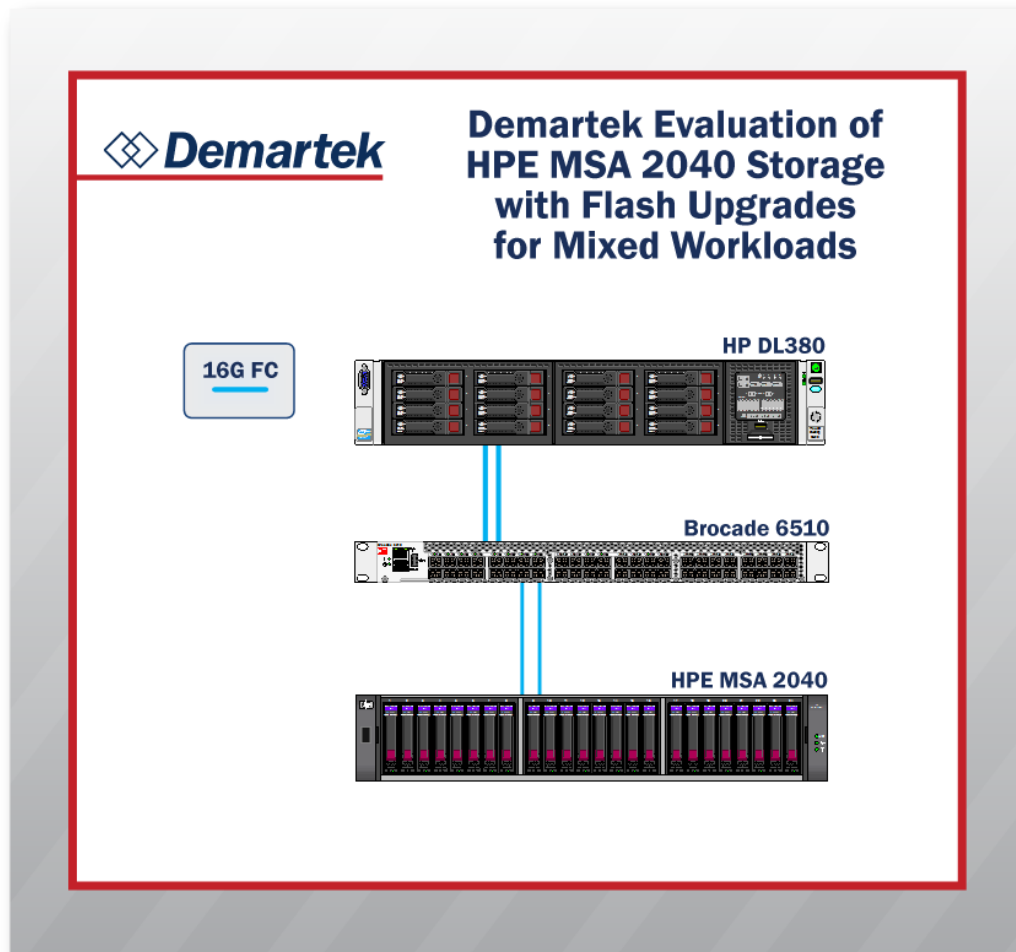


Figure 5 – Test Infrastructure

Server

- ◆ Dual processor rack server
- ◆ 2 Intel E5-5680, 3.33GHz CPUs
- ◆ 144 GB RAM
- ◆ 1 HP SN1000E (16Gb FC) dual port HBA
- ◆ 1 Intel x540-AT2 10Gb NIC
- ◆ vSphere 5.5
- ◆ Microsoft Exchange Jetstress 2013, Microsoft Exchange 2013 ESE files

VMs

- ◆ Active Directory
 - Windows 2012 R2, 8GB RAM, 4 cores
- ◆ Exchange Jetstress
 - Windows 2012 R2, 8GB RAM, 8 cores, MS Exchange Jetstress 2013
- ◆ SQL Server
 - Windows 2012 R2, 16GB RAM, 8 cores, SQL Server 2012 SP2, MS Benchcraft
- ◆ SMB Fileserver
 - Windows 2012 R2, 6GB RAM, 2 cores
- ◆ SharePoint
 - Windows 2012 R2, 8GB RAM 8 cores, SharePoint 2013
- ◆ Webserver
 - Ubuntu 14.0.1, 2GB RAM, 2 cores, NIGNX Webserver
- ◆ vCenter Server
 - Windows 2008 R2, 8GB RAM, 4 cores

Fibre Channel Switch

- ◆ Brocade 6510 16Gb Fibre Channel Switch

Storage Array

- ◆ HPE MSA 2040 Storage array – Utilizing standard Advanced Virtualization features in Base array.
- ◆ Read Cache is a standard feature. MSA Performance Tiering license (optional) was used in SSD Tiering scenarios.
- ◆ 20 900GB 10k RPM 6Gb SAS HDD
 - ◇ 10 drive RAID 6 per storage controller
 - ◇ 550 GB volume for VM files datastore on controller 1
 - ◇ 5000 GB volume for application datastore on controller 1
 - ◇ 5200 GB volume for application datastore on controller 1
- ◆ 4 400GB SSD
 - ◇ 1 drive per storage controller for read caching
 - ◇ 2 drives RAID 1 per storage controller for Performance Tier
- ◆ 4 16Gb FC target ports per controller (1 port active per controller)

Appendix B – Array Volume Configuration

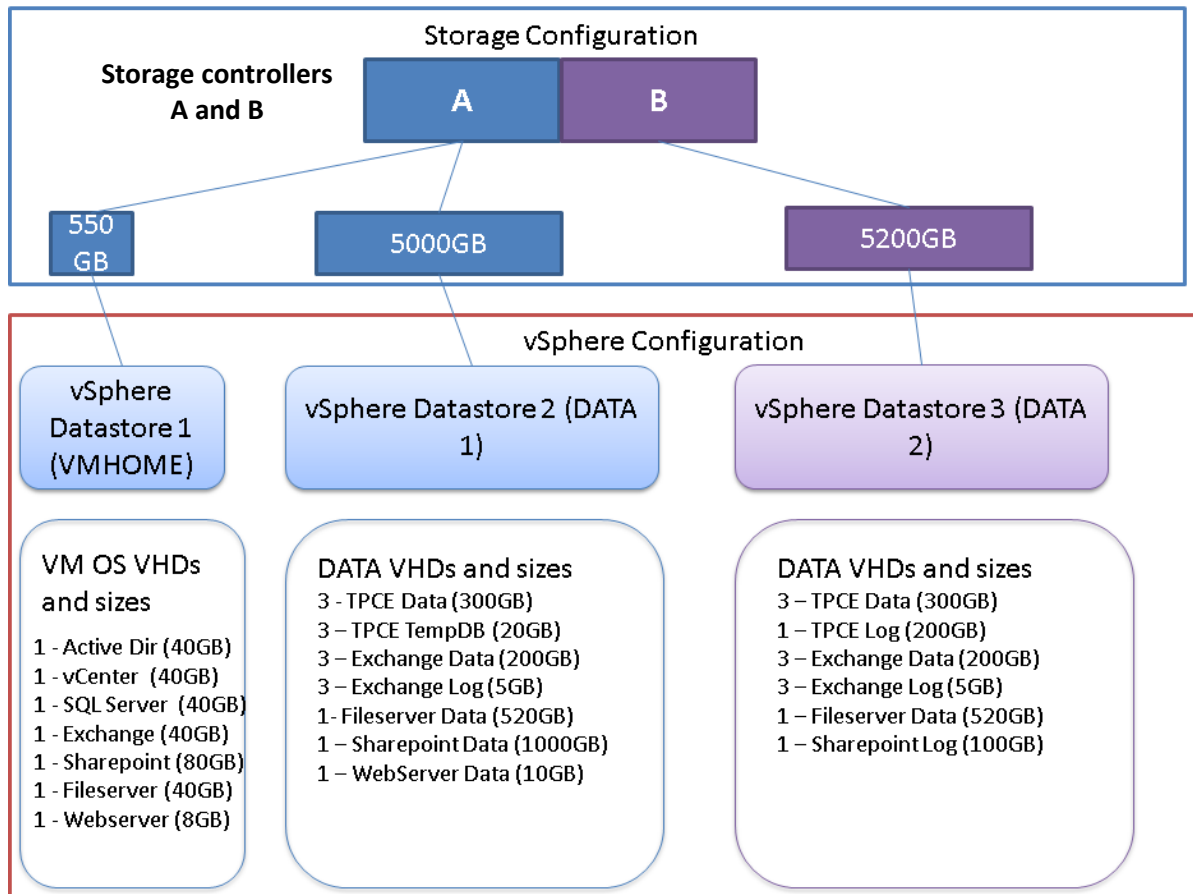


Figure 6 - Array and vSphere volume and datastore configuration

Appendix C – Workload definitions

Exchange Workload

We installed Microsoft Exchange Jetstress 2013 to simulate an Exchange server. While not a true mail server, Jetstress models the I/O loading and patterns that would be generated by an Exchange 2013 server, including email database reads and writes as well as logging. Its configuration includes the number and size of mailboxes, number of mailbox databases and logs, and intensity of use, among other things. Jetstress uses the same Extensible Storage Engine files as Microsoft Exchange to ensure that Jetstress modelled performance is representative of a bona fide Exchange Server of the same version. We configured 250 mailboxes and 6 mailbox databases with a single replica each. Every mailbox was allowed up to 2 GB of space. The workload was configured to use all the available drive space equally, which can be cache unfriendly.

Jetstress was used as the primary benchmark by which we determined whether the storage was able to support our multiple workload environment. If Jetstress could not certify the storage for an average mail server load (.2 IOPS per mailbox for Exchange 2013) we would declare the multiple workloads to be too much for the storage array. Jetstress certifies that an environment is able to support a configuration when average database latencies are below 20 milliseconds.

Demartek also tested the MSA 2040 Storage for suitability as a standalone mail server in another report that can be viewed on the Demartek website.²

SQL Workload

SQL Server 2012 SP2 was installed to support a transactional database workload. This OLTP workload performs real transactions that might be executed by database application users as well as background transactions from automated processes. The workload models a financial brokerage firm with customers who generate transactions related to trades, account inquiries, and market research. The brokerage firm in turn interacts with financial markets to execute orders on behalf of the customers and updates relevant account information. This workload consists of a mixture of mostly reads with some writes to its database.

² Demartek, “Demartek Evaluation – Microsoft Exchange for the Small-to-Medium Business on HPE MSA 2040 Storage”, http://www.demartek.com/Demartek_HPE_MSA_2040_Exchange_Server_Environment_2016-02.html

The database was built with 1TB of data and logs. For our hypothetical business, we assumed that 10 users might be simultaneously accessing this application at any one time. It was not our intention to benchmark a database system. The data gathered by this testing is significant to this evaluation only and cannot be compared with other database performance results.

For more information on running transactional SQL workloads on the HPE MSA 2040 Storage in a standalone capacity please view the Demartek whitepaper “Improving SQL Server Database Performance with HPE MSA 2040 Storage Flash Upgrades”³.

Microsoft SharePoint Workload

Microsoft SharePoint is a common workspace collaboration tool, accessed through a browser from a variety of client devices. We assumed a sixth of our hypothetical company, 25 users, would be simultaneously using SharePoint. With Neotys NeoLoad web performance testing tools, we created a workload that included general navigation of the collaboration space plus uploads and downloads of files. Many of the users were assumed to be accessing and updating related documents, as project teams would do in real-life.

Web Workload

Using the NGINX open source webserver software and an Ubuntu server virtual machine, we deployed a read-only webserver. 25 client device connections were maintained at all times, hitting pages in a semi-random fashion again using the Neotys NeoLoad web performance testing to drive the demand.

Fileserver workload

A basic Windows SMB network share of 1TB was used to create a network fileshare. A server outside of the test environment executed an Iometer script against the network drive. 12 workers ran Iometer access specifications performing 4KB and 8KB random and sequential I/O. A 10 Gb/s Ethernet connection was deployed between the client and fileserver VM. We simulated “hotspots” where several users would perform the same I/O

³ Demartek, “Improving SQL Server Database Performance with HPE MSA 2040 Storage Flash Upgrades”, http://www.demartek.com/Demartek_HPE_MSA_2040_SQL_Server_Environment_2016-02.html



across the same limited portion of the virtual drive, mimicking how certain files tend to be accessed more regularly than others in normal operations.

Appendix D – Fileserver Iometer test file

```

Version 1.1.0
TEST SETUP
=====
Test Description

Run Time
'      hours  minutes  seconds
      24      10      0
Ramp Up Time (s)
0
Default Disk Workers to Spawn
NUMBER_OF_CPUS
Default Network Workers to Spawn
0
Record Results
ALL
Worker Cycling
'      start  step  step type
      1      1  LINEAR
Disk Cycling
'      start  step  step type
      1      1  LINEAR
Queue Depth Cycling
'      start  end  step  step type
      1      32  2  EXPONENTIAL
Test Type
NORMAL
END test setup
RESULTS DISPLAY
=====
Record Last Update Results,Update Frequency,Update Type
DISABLED,1,WHOLE_TEST
Bar chart 1 statistic
Total I/Os per Second
Bar chart 2 statistic
Total MBs per Second (Decimal)
Bar chart 3 statistic
Average I/O Response Time (ms)
Bar chart 4 statistic
Maximum I/O Response Time (ms)
Bar chart 5 statistic
% CPU Utilization (total)
Bar chart 6 statistic
Total Error Count
END results display
ACCESS SPECIFICATIONS
=====
'Access specification name,default assignment
4 KiB; 100% Read; 25% random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
4096,100,100,25,0,1,0,0
'Access specification name,default assignment
4 KiB; 10% Read; 0% random,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
4096,100,10,0,0,1,0,0
'Access specification name,default assignment
8 KiB; mixed; hotspots,NONE
'size,% of size,% reads,% random,delay,burst,align,reply
8192,10,100,10,0,1,0,0

```

8192,10,0,10,0,1,0,0
8192,20,10,0,0,1,0,0
8192,60,100,10,0,1,0,0

'END access specifications
'MANAGER LIST

=====
'Manager ID, manager name
1,WIN-08GAJTULK4J
'Manager network address

'Worker

Worker 1

'Worker type
DISK

'Default target settings for worker

'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
1,DISABLED,1,DISABLED,0

'Disk maximum size,starting sector,Data pattern
0,0,1

'End default target settings for worker

'Assigned access specs
4 KiB; 100% Read; 25% random

'End assigned access specs

'Target assignments

'Target
F:\\10.0.9.104\

'Target type
DISK

'End target

'End target assignments

'End worker

'Worker

Worker 2

'Worker type
DISK

'Default target settings for worker

'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
1,DISABLED,1,DISABLED,0

'Disk maximum size,starting sector,Data pattern
0,0,1

'End default target settings for worker

'Assigned access specs
4 KiB; 100% Read; 25% random

'End assigned access specs

'Target assignments

'Target
F:\\10.0.9.104\

'Target type
DISK

'End target

'End target assignments

'End worker

'Worker

Worker 3

'Worker type
DISK

'Default target settings for worker

'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
1,DISABLED,1,DISABLED,0

'Disk maximum size,starting sector,Data pattern
0,0,1

```
'End default target settings for worker
'Assigned access specs
    4 KiB; 100% Read; 25% random
'End assigned access specs
'Target assignments
'Target
    F:\\10.0.9.104\
'Target type
    DISK
'End target
'End target assignments
'End worker
'Worker
    Worker 4
'Worker type
    DISK
'Default target settings for worker
'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
    1,DISABLED,1,DISABLED,0
'Disk maximum size,starting sector,Data pattern
    0,0,1
'End default target settings for worker
'Assigned access specs
    4 KiB; 100% Read; 25% random
'End assigned access specs
'Target assignments
'Target
    F:\\10.0.9.104\
'Target type
    DISK
'End target
'End target assignments
'End worker
'Worker
    Worker 5
'Worker type
    DISK
'Default target settings for worker
'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
    1,DISABLED,1,DISABLED,0
'Disk maximum size,starting sector,Data pattern
    0,0,1
'End default target settings for worker
'Assigned access specs
    4 KiB; 10% Read; 0% random
'End assigned access specs
'Target assignments
'Target
    F:\\10.0.9.104\
'Target type
    DISK
'End target
'End target assignments
'End worker
'Worker
    Worker 6
'Worker type
    DISK
'Default target settings for worker
'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
    1,DISABLED,1,DISABLED,0
```

```
'Disk maximum size,starting sector,Data pattern
    0,0,1
'End default target settings for worker
'Assigned access specs
    4 KiB; 10% Read; 0% random
'End assigned access specs
'Target assignments
'Target
    F:\\10.0.9.104\f
'Target type
    DISK
'End target
'End target assignments
'End worker
'Worker
    Worker 7
'Worker type
    DISK
'Default target settings for worker
'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
    1,DISABLED,1,DISABLED,0
'Disk maximum size,starting sector,Data pattern
    0,0,1
'End default target settings for worker
'Assigned access specs
    4 KiB; 10% Read; 0% random
'End assigned access specs
'Target assignments
'Target
    F:\\10.0.9.104\f
'Target type
    DISK
'End target
'End target assignments
'End worker
'Worker
    Worker 8
'Worker type
    DISK
'Default target settings for worker
'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
    1,DISABLED,1,DISABLED,0
'Disk maximum size,starting sector,Data pattern
    0,0,1
'End default target settings for worker
'Assigned access specs
    4 KiB; 10% Read; 0% random
'End assigned access specs
'Target assignments
'Target
    F:\\10.0.9.104\f
'Target type
    DISK
'End target
'End target assignments
'End worker
'Worker
    Worker 9
'Worker type
    DISK
'Default target settings for worker
```

```
'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
  1,DISABLED,1,DISABLED,0
'Disk maximum size,starting sector,Data pattern
  10737418240,0,1
'End default target settings for worker
'Assigned access specs
  8 KiB; mixed; hotspots
'End assigned access specs
'Target assignments
'Target
  F:\\10.0.9.104\
'Target type
  DISK
'End target
'End target assignments
'End worker
'Worker
  Worker 10
'Worker type
  DISK
'Default target settings for worker
'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
  1,DISABLED,1,DISABLED,0
'Disk maximum size,starting sector,Data pattern
  10737418240,0,1
'End default target settings for worker
'Assigned access specs
  8 KiB; mixed; hotspots
'End assigned access specs
'Target assignments
'Target
  F:\\10.0.9.104\
'Target type
  DISK
'End target
'End target assignments
'End worker
'Worker
  Worker 11
'Worker type
  DISK
'Default target settings for worker
'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
  1,DISABLED,1,DISABLED,0
'Disk maximum size,starting sector,Data pattern
  10737418240,0,1
'End default target settings for worker
'Assigned access specs
  8 KiB; mixed; hotspots
'End assigned access specs
'Target assignments
'Target
  F:\\10.0.9.104\
'Target type
  DISK
'End target
'End target assignments
'End worker
'Worker
  Worker 12
'Worker type
```

```
DISK
'Default target settings for worker
'Number of outstanding IOs,test connection rate,transactions per connection,use fixed seed,fixed seed value
    1,DISABLED,1,DISABLED,0
'Disk maximum size,starting sector,Data pattern
    10737418240,0,0
'End default target settings for worker
'Assigned access specs
    8 KiB; mixed; hotspots
'End assigned access specs
'Target assignments
'Target
    F:\\10.0.9.104\
'Target type
    DISK
'End target
'End target assignments
'End worker
'End manager
'END manager list
Version 1.1.0
```

The original version of this document is available at:
http://www.demartek.com/Demartek_HPE_MSA_2040_Mixed_Workload_Environment_2016-02.html on the Demartek website.

Intel and Xeon are registered trademarks of Intel Corporation.

Neotys and NeoLoad are registered trademarks of Neotys Corporation

Microsoft, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Demartek is a trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.