Demartek

Dell EMC vSAN Mixed Enterprise Workloads with Toshiba 12Gb/s SAS SSDs

Mixed Enterprise Workloads with Dell EMC PowerEdge R730 servers, VMware vSAN, and Toshiba PX05S Series 12Gb/s SAS SSDs

Executive Summary

Hyper-converged infrastructure (HCI) is a newer technology where a cluster of hypervisor server nodes are deployed in a single chassis containing tightly-integrated server, networking, and storage technology. VMware vSAN manages HCI distributed storage and is part of the VMware vSphere Hypervisor. Dell EMC provides VMware vSAN Ready Nodes that combine a wide range of Dell hardware and VMware vSAN software into a ready-to-order package, with Toshiba PX05S Series 12Gb/s SAS SSDs. These Dell EMC vSAN Ready Nodes are validated and configured to meet your hyper-converged workload needs, including value-optimized configurations for smaller projects, storage-dense configurations that require large storage capacities or compute-dense solutions for compute-intensive workloads. Dell certifies its vSAN ready node on 2U dual processor Dell EMC R730 servers which are suitable for multiple application workloads.

Demartek deployed a three-node all-flash Dell EMC vSAN cluster with Dell PowerEdge R730 servers. The performance of this cluster should be the same as that of a cluster of Dell EMC vSAN ready nodes. Each server had 5 Toshiba PX05S Series 3.84 TB 12Gb/s SAS SSDs and was running VMware ESXi 6.5.

Across the virtual machines (VMs) running on this cluster, several enterprise workloads were deployed including:

> VMware vCenter

> Windows Server Active Directory

> 3x Microsoft SQL Server (DVDStore2)

> Microsoft Exchange Jetstress (4000 mailboxes with heavy I/O profile)

> File server and clients

> Web server simulation

Key Findings

> With Toshiba PX05S Series 12Gb/s SAS SSDs, for the combined workloads, we achieved an average read latency of approximately 500 microseconds (500 µs or ½ millisecond).

> With Toshiba PX05S Series 12Gb/s SAS SSDs, for the combined workloads, we achieved an average write latency of approximately 2 milliseconds (2 ms).

Dell EMC vSAN Mixed Enterprise Workloads with Toshiba 12Gb/s SAS SSDs

Demartek
September 2017

Server Hardware

Three Dell EMC PowerEdge R730 servers were used for the vSAN cluster. Each server included:

> 2x Intel® Xeon® E5-2698 v4 processors, 2.2 GHz, 40 total cores, 80 total threads

> 512 GB memory

> VMware ESXi 6.5

> 5x Toshiba PX05S Series 3.84 TB 12Gb/s SAS SSDs

> Intel x520 10GbE NIC

Storage Hardware

Five Toshiba PX05S Series 3.84TB 12Gb/s SAS SSDs were used for data storage in each node of the vSAN cluster. 7.3TB was allocated to guest VMs running application workloads. The R730 servers use an internal SD card as the boot drive.

Switches

The cluster network was connected via a Dell EMC S4048-ON 10GbE switch.

Configuration

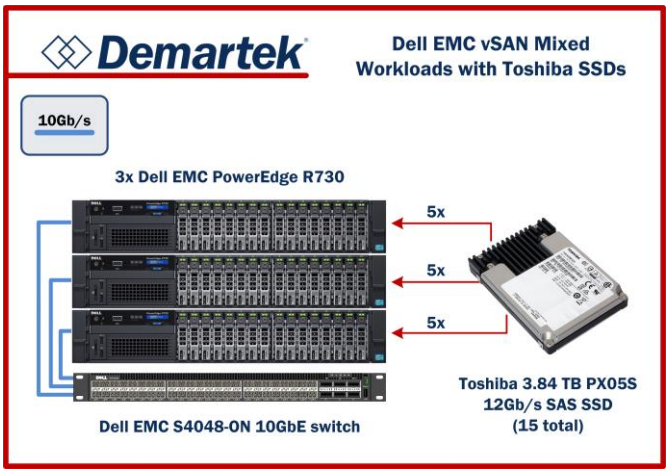

Applications and Virtual Machines

We created 9 virtual machines. Application workloads chosen to simulate the type of work often done in an enterprise datacenter were distributed across the nodes in multiple virtual machines as listed below:

> Win2016-AD (Active Directory)

> vCenter Server Appliance

> Win2016-Exch (Jetstress Exchange)

> Win2016-FS (File Server)

> Win2016-FSClient (Fileshare client)

> Win2016-WS (Webserver)

> Win2016-SQL2016 (DVDStore2)

> Win2016-SQL2016-DB2 (DVDStore2)

> Win2016-SQL2016-DB3 (DVDStore2)

Performance Results

Processor and Memory Utilization

Memory utilization is a measure of the amount of memory used by the applications running. The average memory utilization across the cluster was just under 60GB. Processor or CPU utilization is a measure of the amount of CPU processing that occurs as workloads run on a system. Some workloads are more CPU intensive than others, while some are more I/O intensive. The average peak workload on each node was approximately 11% CPU utilization.

The low processor and memory utilization rates indicate that the Dell PowerEdge R730 server cluster was able to support all workloads efficiently with plenty of room left for future growth.

Dell EMC vSAN Mixed Enterprise Workloads with Toshiba 12Gb/s SAS SSDs

Read and Write Bandwidth

Bandwidth or throughput is a measure of the amount of data transferred to or from a storage device or set of storage devices, per second. This is typically measured in kilobytes per second (KBPS), megabytes per second (MBPS) or gigabytes per second (GBPS). Storage performance metrics are measured in bytes per second. By contrast network performance is measured in bits per second.

Storage performance also frequently separately identifies reads and writes. Reads are requests from hosts to retrieve data from storage and writes are requests by hosts to send data to storage. Many workloads are asymmetrical with respect to reads and writes – they retrieve different amounts of data from storage than they send to storage.

Looking at our graphs, we can see that our mixed workload has more read bandwidth than write bandwidth.

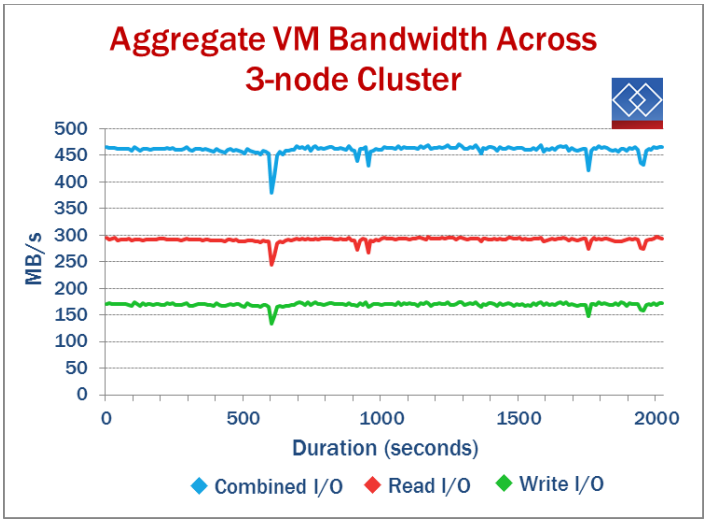

IOPs

Another common storage performance metric is reads and writes per second, or I/Os per second (IOPS). This is a measure of the number of I/O requests issued by a host server per second.

Workloads that are transaction oriented tend to have higher IOPS rates than bulk transfer types of applications such as backup or video streaming. As is typical for real-world workloads, IOPS rates are not flat but fluctuate, sometimes with occasional bursts or valleys. In our testing the workloads produce a mostly steady load, with an average of almost 35,000 IOPS.

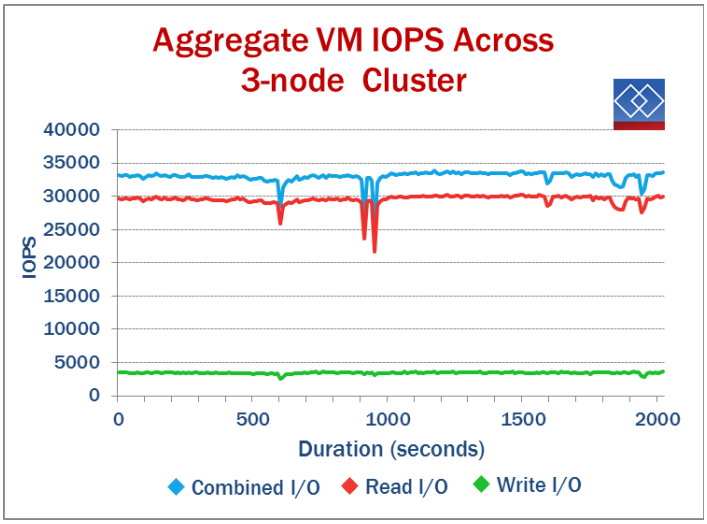

Notice that when we compare write I/O to read I/O, the delta is much larger than when we compared write bandwidth to read bandwidth. This suggests that while there were less write requests, that on average each individual write request was larger, or had a larger block size, than each read request. This is important to remember when we consider response time.

I/O Response Time

Response time is a measure of the time taken to complete an I/O request; it is also known as latency. Latency is frequently measured in milliseconds (ms), which is thousandths of a second.

Running a set of mixed enterprise workloads in an all-flash vSAN hyperconverged environment powered by Dell EMC vSAN and Toshiba SAS SSDs delivered a multi-workload, network storage solution with response times of 2ms and less. The average read latency was approximately 500 microseconds and the average write latency was approximately 2 milliseconds.

Dell EMC vSAN Mixed Enterprise Workloads with Toshiba 12Gb/s SAS SSDs

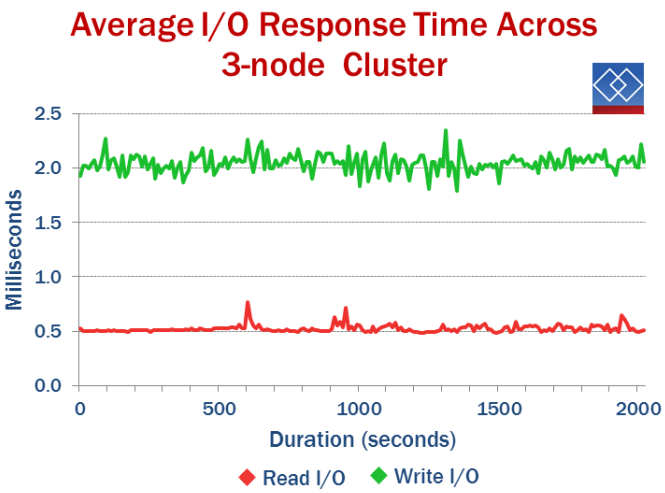

Notice that write response times are higher than read response times. We noted earlier that in this workload writes had a larger block size on average than reads, which explains the higher latency. This type of latency and large block size is typical of writes to log files. The workload application queues a large amount of data and writes it all at once to a log; the large amount of data makes the write take longer to complete. As such, the higher write latency reflects the character of our mixed workload more than it reflects the storage used. That we were able to get just 2ms of write latency with the larger block sizes found in this workload is an impressive accomplishment for our storage.

Summary and Conclusion

Hyperconverged infrastructure is becoming more popular as a way to pack more compute and storage into less rack units, and the Dell EMC vSAN with Toshiba SSDs is an excellent server platform for HCI. We ran a set of mixed workloads in nine virtual machines spread across three nodes of a Dell PowerEdge R730 vSAN server cluster. These were enterprise workloads, including VMware vCenter, Windows Active Directory, Microsoft Exchange Jetstress, Fileserver, Webserver, and DVDStore2 Microsoft SQL workloads. The single Dell PowerEdge R730 vSAN cluster with the help of the Toshiba PX05S Series 3.84TB 12Gb/s SAS drives was able to handle all these workloads, providing low latency, high bandwidth and leaving plenty of processor for other applications.

Dell PowerEdge R730 All-flash vSAN cluster nodes are the way to go when performance is key and Toshiba PX05S Series 3.84TB 12Gb/s SAS drives deliver a large amount of flash in a small package, ideal for busy virtualized environments. The Dell PowerEdge R730 server supports enough CPU and memory to satisfy mixed enterprise workloads, and when used in vSAN clusters with high-speed networking infrastructure, makes a powerful VMware vSAN platform.

The most current version of this report is available at www.demartek.com/Dell-vSAN-Toshiba on the Demartek website.