The Performance of SAS

SAS provides a strong performance for a variety of enterprise application workloads.

Executive Summary

Serial attached SCSI, or SAS provides a reliable enterprise-grade interface designed specifically for computer storage. While well-established, SAS technology continues to adapt to the address the demanding requirements of new technologies in today's data centers. SAS offers a wide and robust ecosystem including host bus adapters (HBAs), RAID adapters, expanders, cables, connectors, and system backplanes. SAS storage systems can scale all the way up to large external storage systems with hundreds or thousands of storage devices.

Server virtualization has been embraced by many organizations, and many are deploying mission-critical applications in virtual machines (VMs). Some have embraced a "virtualization first" policy, which prefers deployment on VMs over physical hardware. The goal of this project is to show the performance and flexibility of a multi-vendor 12Gb/s SAS storage solution running a set of mixed workloads in a real-world, virtualized environment.

For this testing, we configured a multi-vendor collection of 12Gb/s SAS SSDs, a 12Gb/s SAS RAID controller, and a 12GB/s SAS JBOD enclosure in a VMware environment running two different real-world application workloads simultaneously. The 12Gb/s SAS infrastructure performed well, given the complexity of environment, which is representative of a real-world customer environment.

Key Findings

> ***Real-world Latency:*** For the OLTP-only workloads, our real-world virtualized server environment with multiple guest virtual machines achieved read and write latencies below 200 microseconds using a mixture of 12Gb/s SAS infrastructure.

> ***Flexible Configuration:*** The SAS infrastructure and in-box drivers helped to simplify the deployment of a complex storage environment.

Test Configuration – Hardware

Server

> Dell PowerEdge R820

> 4x Intel Xeon E5-4650, 2.7 GHz, 256GB RAM

SAS RAID Controller

> Microsemi / Adaptec 8885Q, 12Gb/s SAS

> 2x internal mini SAS HD (SFF-8643) ports

> 2x external mini SAS HD (SFF-8644) ports

SAS SSDs (12Gb/s)

> 8x HGST HUSMM1616ASS200, 1.6TB each

> 8x Seagate ST1600FM0013, 1.6TB each

> 8x Toshiba PX04SMB080, 800GB each

SAS Storage Enclosure

> Xyratex 24-drive SFF 12Gb/s SAS enclosure

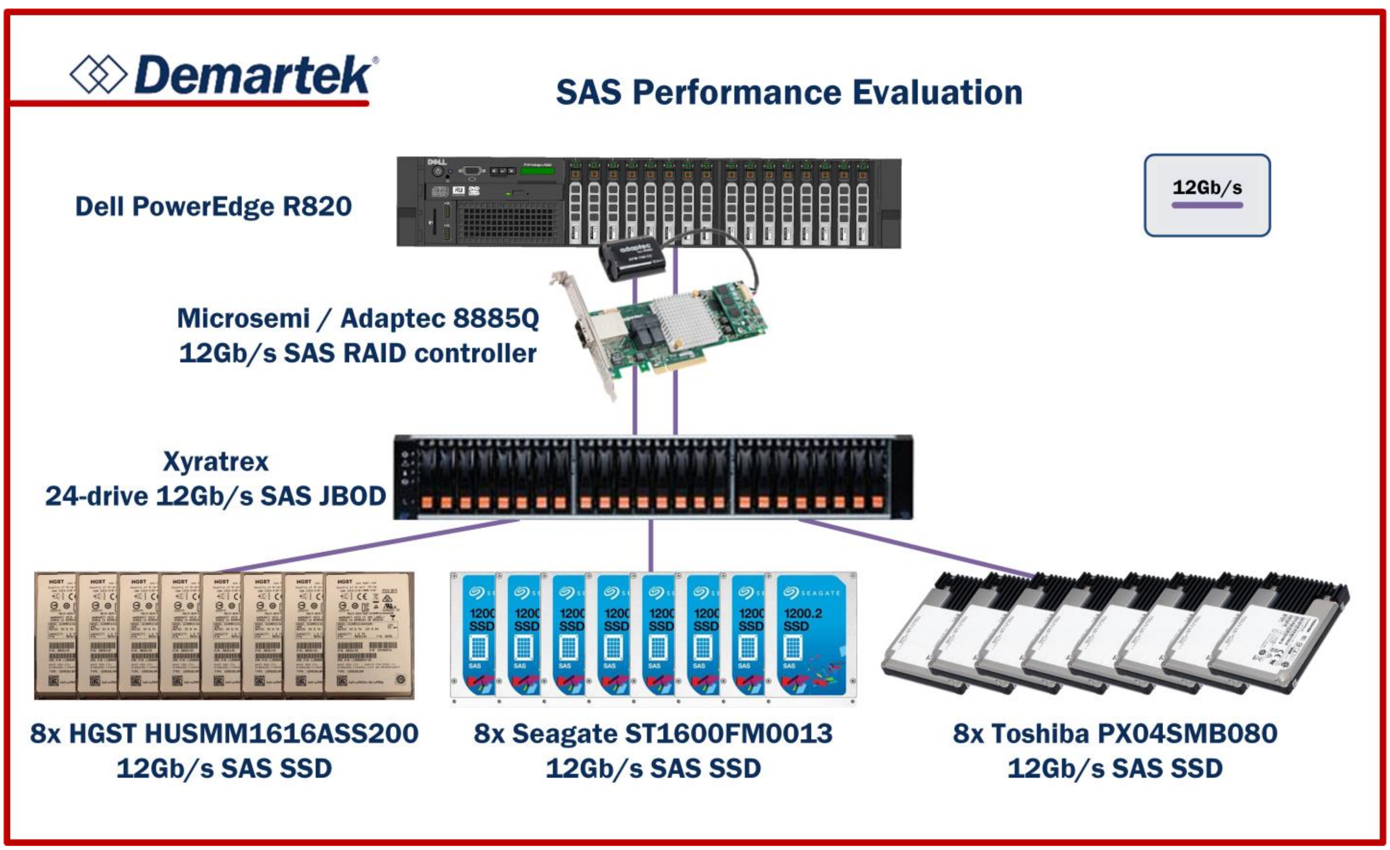

Logical Storage Configuration

The SAS SSDs were logically organized into 800GB volumes that were mirrored (RAID1), using the RAID controller. The first four of these 800GB logical volumes (two mirrors) were used as a 1.46TB LUN for the operating system. The remaining volumes were organized into two large data LUNs of 7.28TB and 5.82TB. The two data LUNs were placed into a VMware datastore.

The diagram below visually depicts the logical storage configuration as presented to VMware.

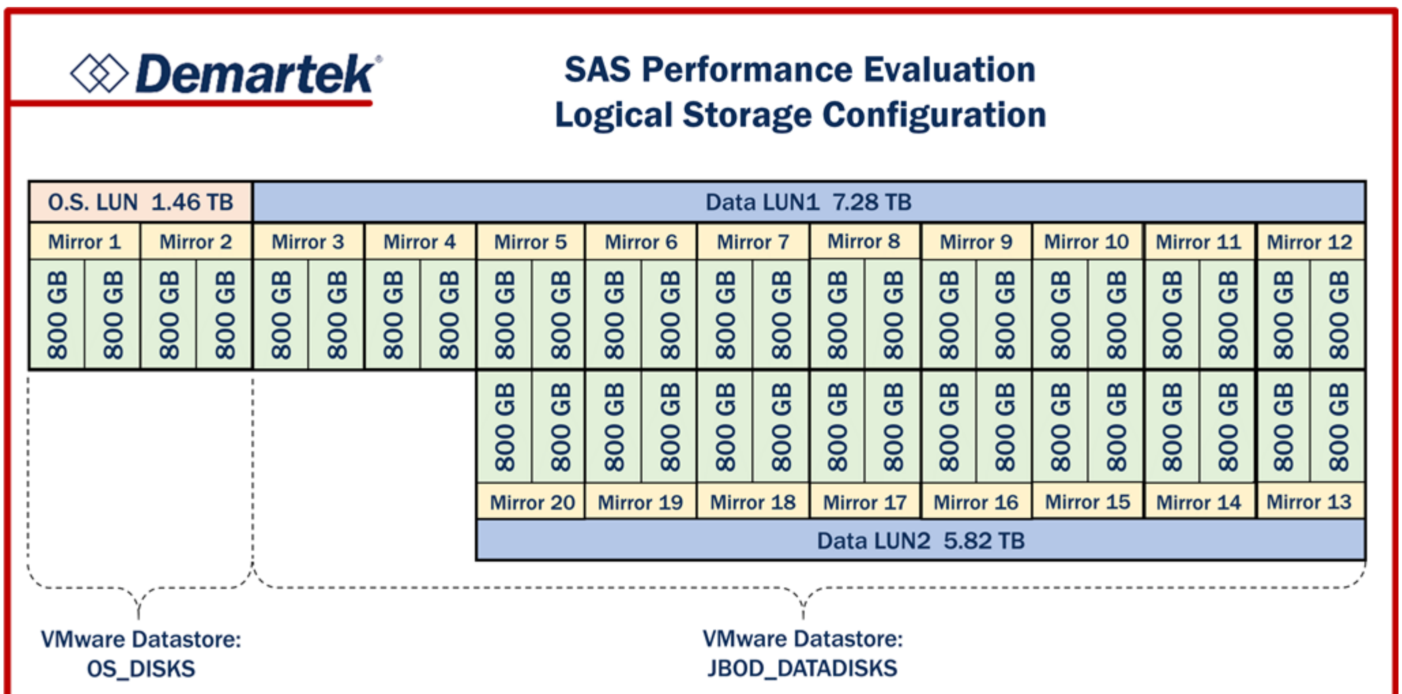

These logical volumes were composed of 8x 800GB SSDs and 16x 1.6TB SSDs partitioned into two logical drives of 800GB each, resulting in a total of 40x 800GB volumes. The RAID1 mirrors were created using the 12Gb/s SAS RAID controller, using drives from the same vendor for each mirror. The LUNs were created as RAID10 stripes. The LUN sizes are formatted capacity.

Real-World Workload Description

An increasing number of organizations are turning to server virtualization for database applications such as Microsoft SQL Server. Virtualizing these workloads allows for allocation of the optimal amount of compute resources to applications, with the ability to easily adjust these resources as requirements change. Deploying these types of applications in virtual machines (VMs) allows workloads to be migrated to different physical servers much more easily than if they were deployed on physical hardware. If needed, these applications can also take advantage of any fault-tolerant capabilities provided by the virtualization layer in addition to fault tolerant abilities of the application, improving disaster planning scenarios.

Two application workloads running in a VMware environment were designed to put a difficult-to-predict workload onto the storage components. The mixture of I/O requests in an environment such as this is often given the name of the “I/O blender” effect. Further, the performance of the storage system is a blend of the performance characteristics of the various types of SSDs used in this configuration.

VMware Environment

A relatively simple VMware environment (vSphere 6.5) was used with two guest VMs each running Windows Server 2016 and Microsoft SQL Server 2016. One guest VM ran the OLTP workload on SQL Server and the other guest ran the data warehousing workload on SQL Server.

OLTP Workload

The On-Line Transaction Processing (OLTP) workload is a variation of an OLTP workload that models a brokerage firm with customers who generate transactions related to trades, account inquiries, and market research. This workload is approximately 90% read and 10% write (not counting the log files that are 100% write). The brokerage firm in turn interacts with financial markets to execute orders on behalf of the customers and updates relevant account information.

The benchmark is “scalable,” meaning that the number of customers defined for the brokerage firm can be varied to represent the workloads of different-size businesses. The benchmark defines the required mix of transactions the benchmark must maintain.

This workload behavior has the following characteristics:

- > This workload emphasizes IOPS (smaller block sizes) and has low, relatively steady latency.
- > This workload typically consumes at least 50% CPU utilization on all or most cores, depending on the test settings.

The customer count was 75,000.

Data Warehousing Workload

This read-intensive (100% read, not counting the log files which are 100% write) workload consists of a suite of business-oriented ad-hoc queries and concurrent data modifications. The queries and the data populating the database have been chosen to have broad industry-wide relevance. This benchmark illustrates decision support systems that examine large volumes of data, execute queries with a high degree of complexity, and give answers to critical business questions.

This workload behavior has the following characteristics:

- > This workload emphasizes throughput (large block sizes) and has variable, sometimes high latency.
- > This workload issues a fixed set of work – faster systems will complete the work in less time.
- > This workload fluctuates and sometimes consumes as much as 100% CPU utilization, depending on the test parameters.

A scale factor of 1000 was used, resulting in a database of approximately 1000GB in size.

Performance Results – Mixed Workload

We collected several metrics from this mixed workload set of tests. Given the complexity of the environment, these are fairly strong performance results.

Block Size

To begin the analysis, the ***I/O block size*** should be noted, which affects all the other metrics. The averages for the workloads are:

- Data warehousing database reads: 279.6KB
- Data warehousing log writes: 1.4KB (infrequent and variable)
- OLTP database reads: 11.8KB
- OLTP database writes: 7.8KB

As is typical for real-world workloads, the I/O block size varies during each run.

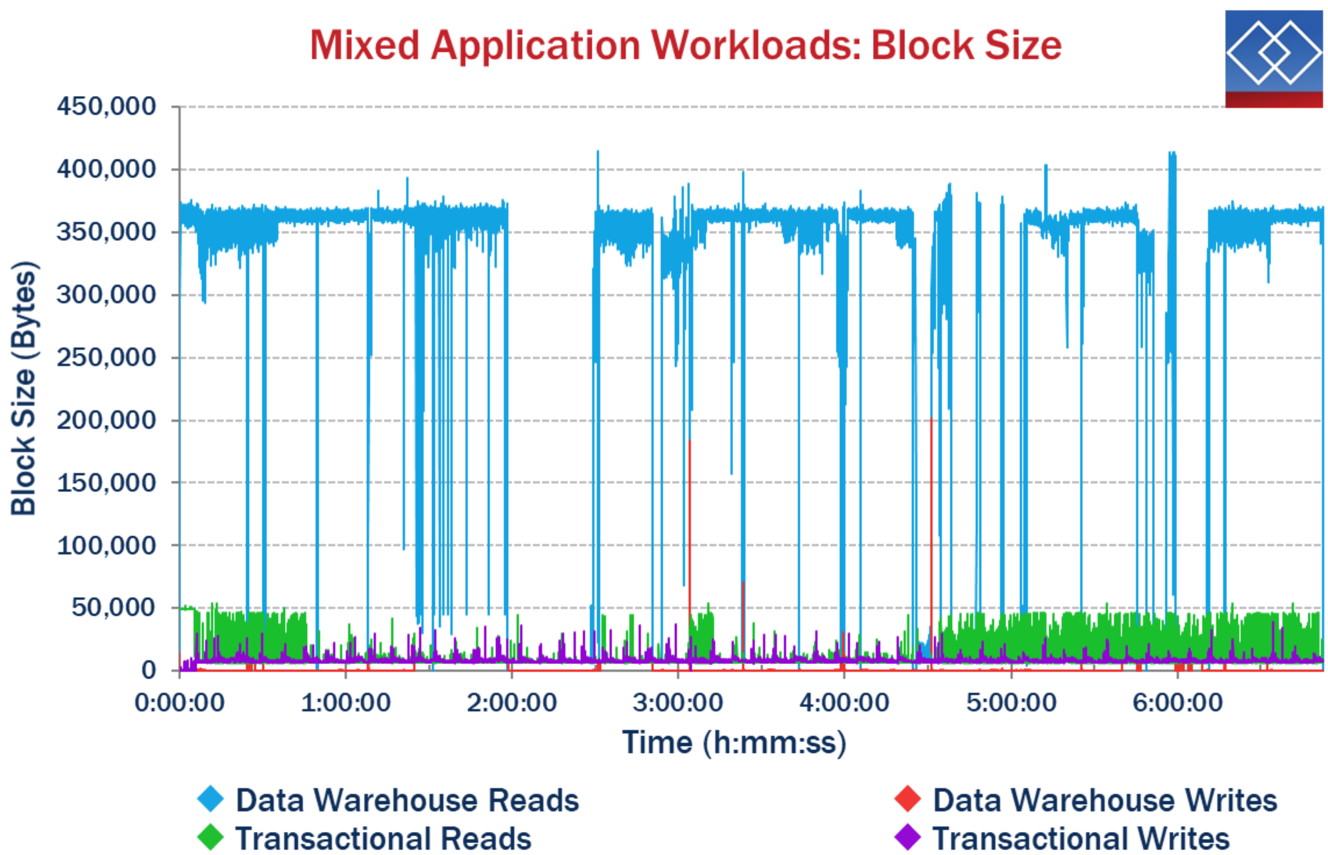

IOPS

Not surprisingly, the OLTP workload achieved higher IOPS than the data warehousing workload, due to its smaller average block size.

Throughput

Conversely, the data warehousing workload achieved much higher throughput, on average, than the OLTP workload. This is expected due to its large I/O block size.

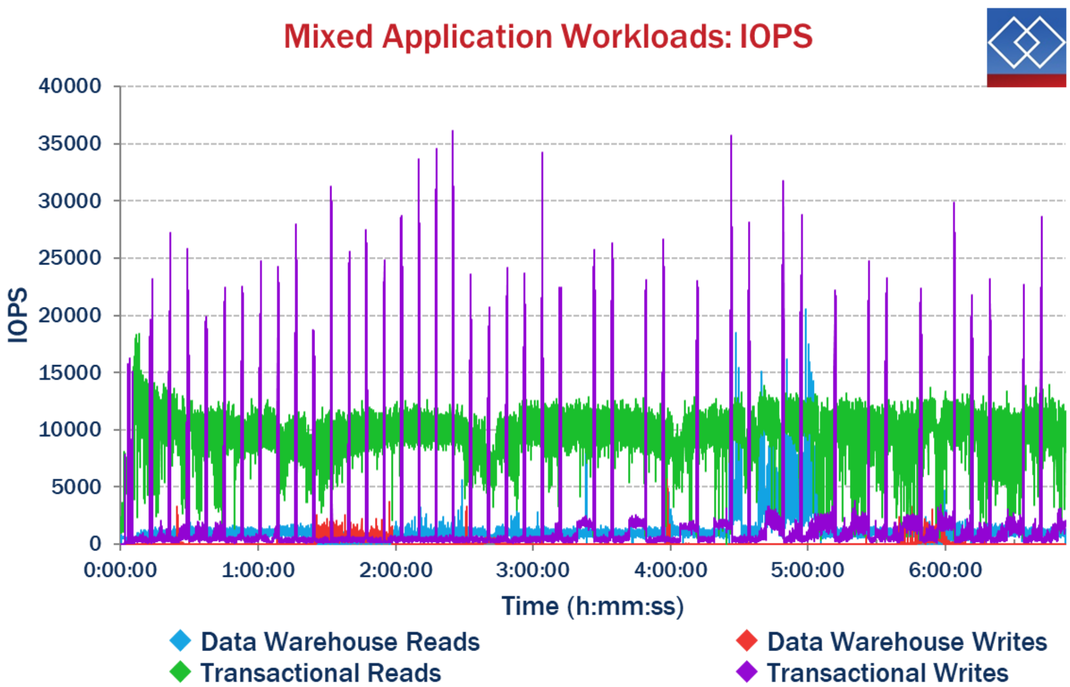

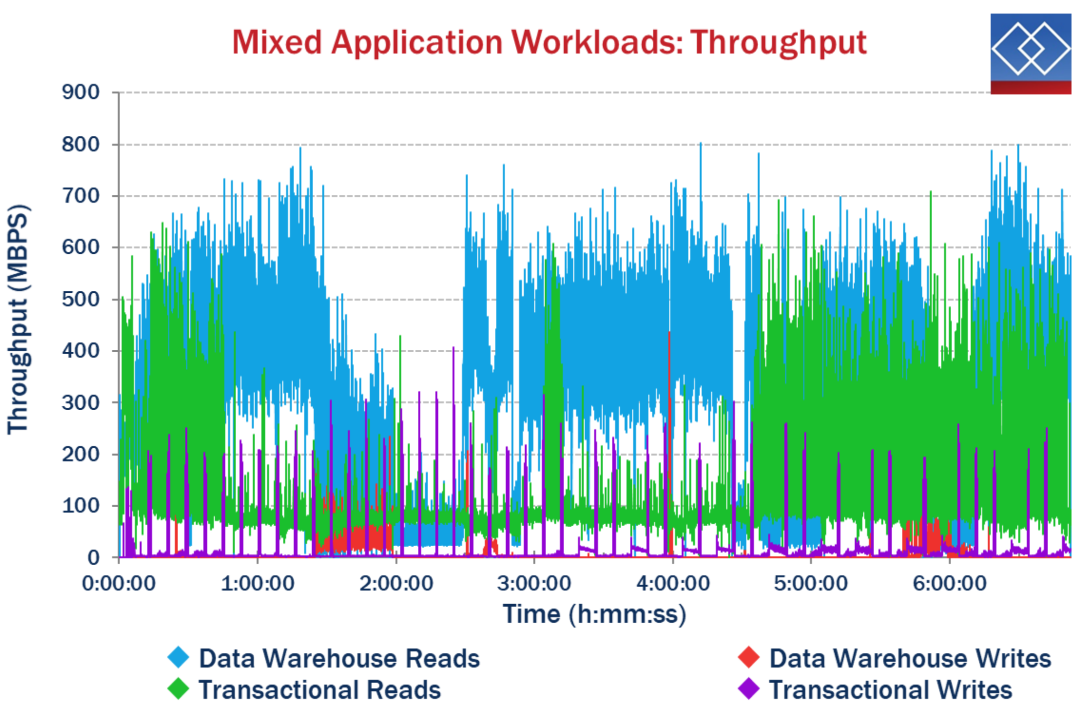

Latency

Not surprisingly, the data warehouse read latencies were higher than the other latencies, again due to the large I/O block size. The latency shown here is as viewed from inside the virtual machine. This latency includes the latency from multiple components or layers, including the SSD, RAID configuration, hypervisor, and guest operating system. In a real-world environment, latency is more than just the drive latency.

In the minimum latencies graph below, we see that the data warehousing read workload negatively impacted the minimum latencies of the other workloads, by raising them above what they would be without the data warehousing workload, especially the OLTP reads (see the OLTP-only section on the next page). This shows that heavy workloads from one VM can affect the storage performance of workloads running on different VMs.

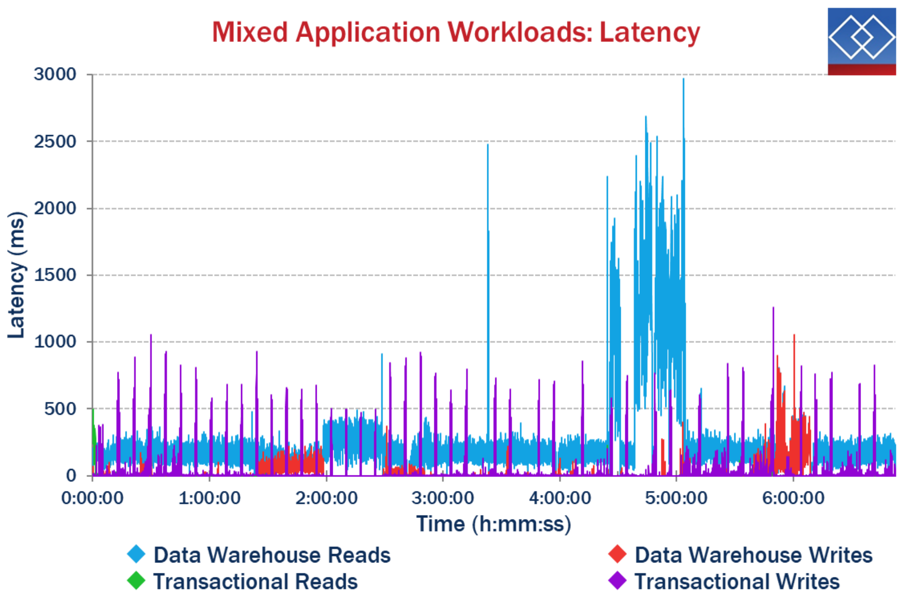

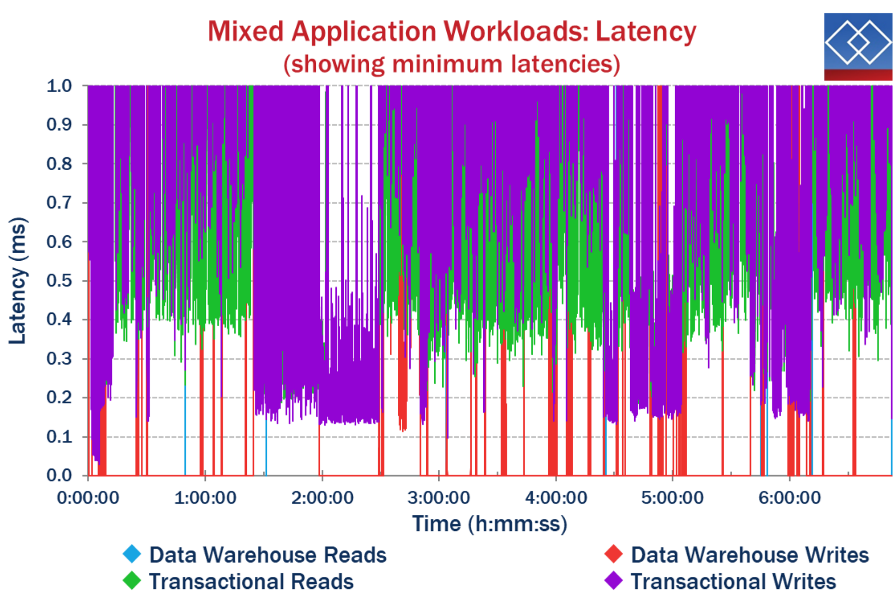

Memory Utilization

We found that the usage of system memory was fairly constant during these runs.

CPU Utilization

Processor utilization is directly related to the type of work the application performs. In a real-world application environment, the CPU is performing calculations, loading data from and storing data into memory, and performing I/O operations. As is typical, the OLTP workload consumes a relatively consistent CPU utilization (~50%) while the data warehousing workload generally consumes lower CPU resources but is a bit more variable.

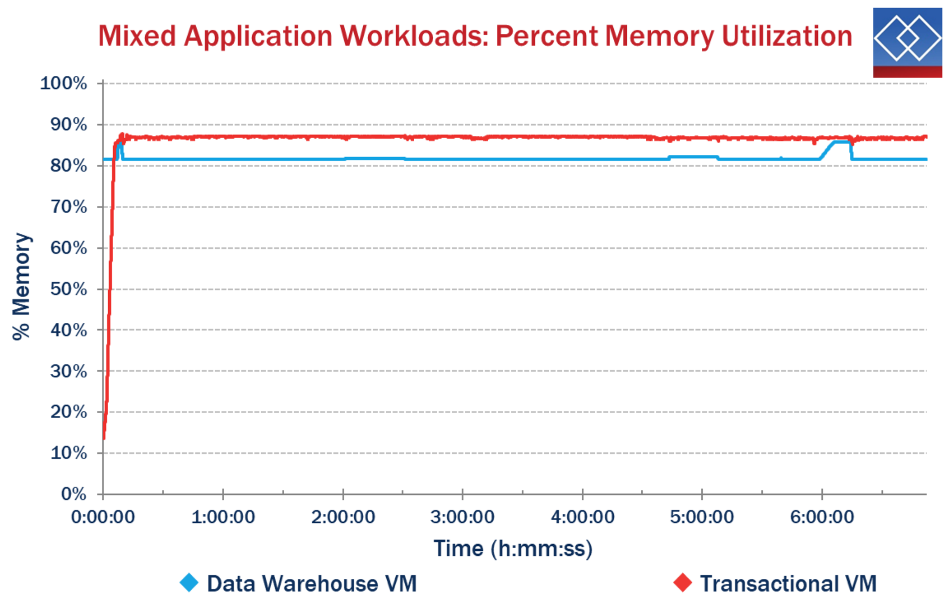

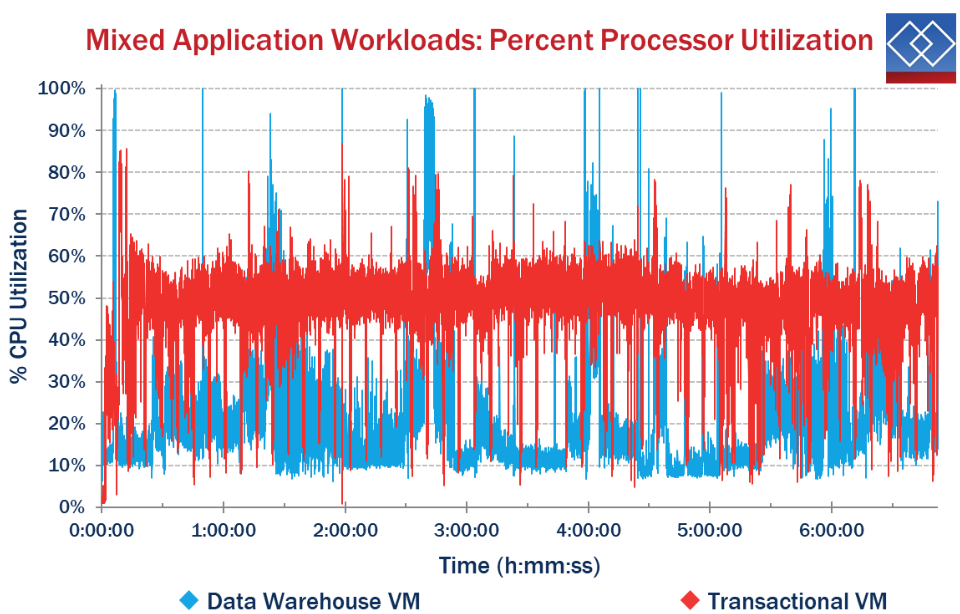

Performance Results – OLTP Only

Running the OLTP workload separately gives us some insight into this workload that we might not notice when it is mixed with other workloads.

Block Size

A large number of the block sizes are 8KB (8192 bytes) but they are not limited to that block size. Some reads and writes spike much higher. This is the nature of real-world workloads.

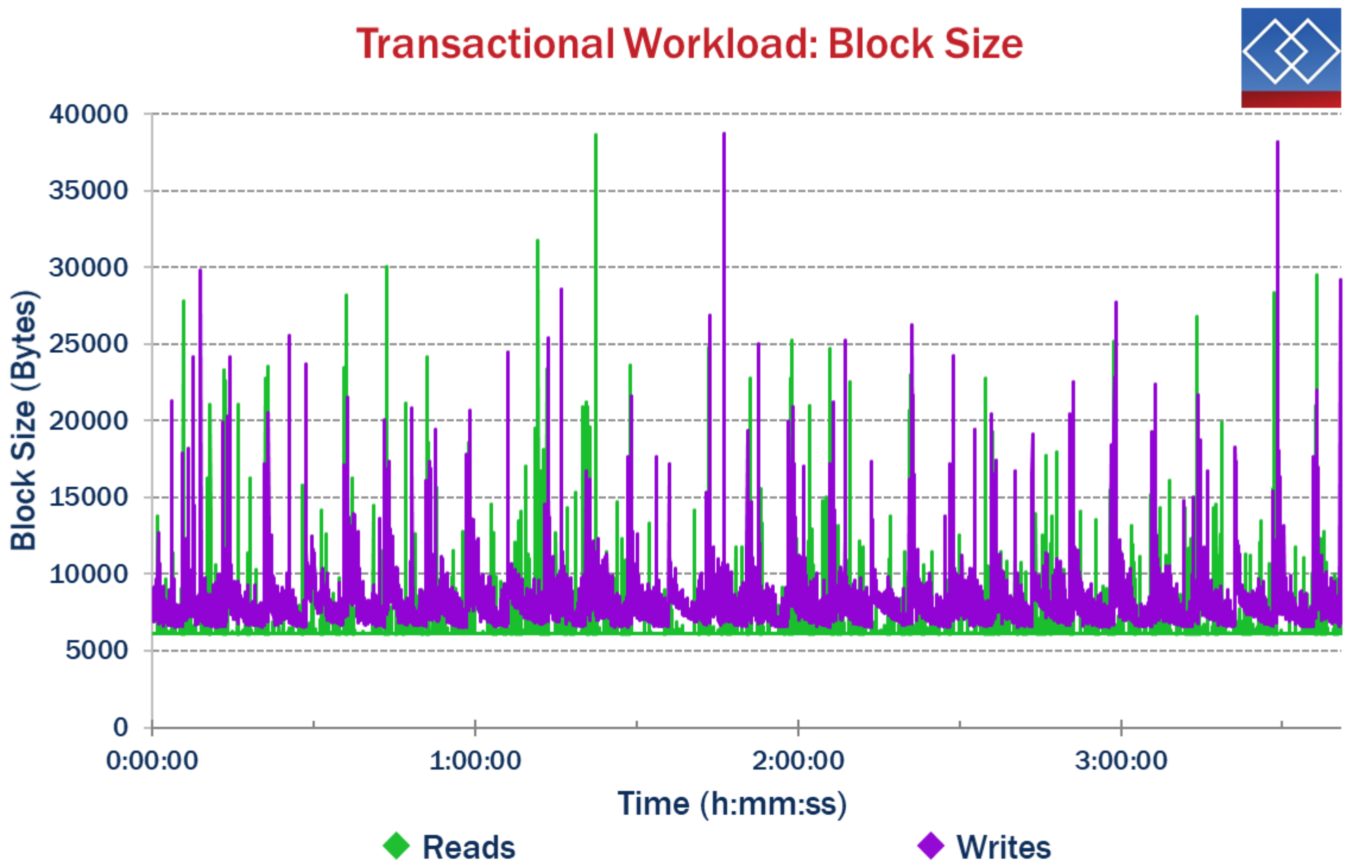

IOPS

The read IOPS are fairly steady for the OLTP workload, but the write IOPS depend on when transaction records are added or updated and log writes take place. The writes correlate to the rate at which updates are made.

Throughput

Throughput for the OLTP workload is not very high. The read activity has a relatively constant base but does spike occasionally. Writes follow a somewhat predictable pattern, at least with respect the number of updates that occur, and indirectly to elapsed time.

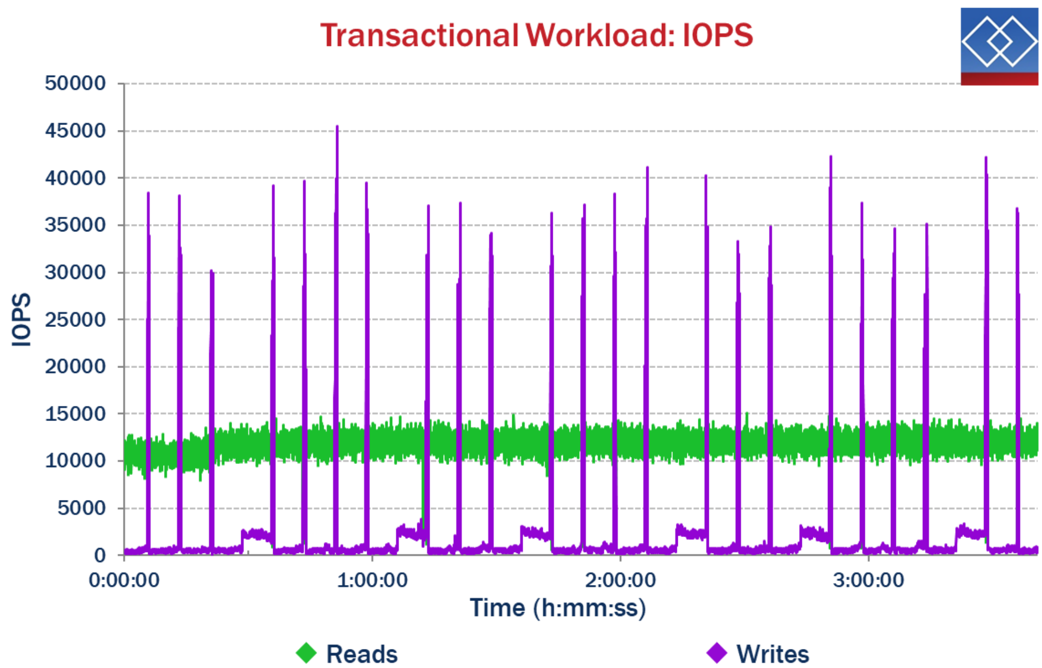

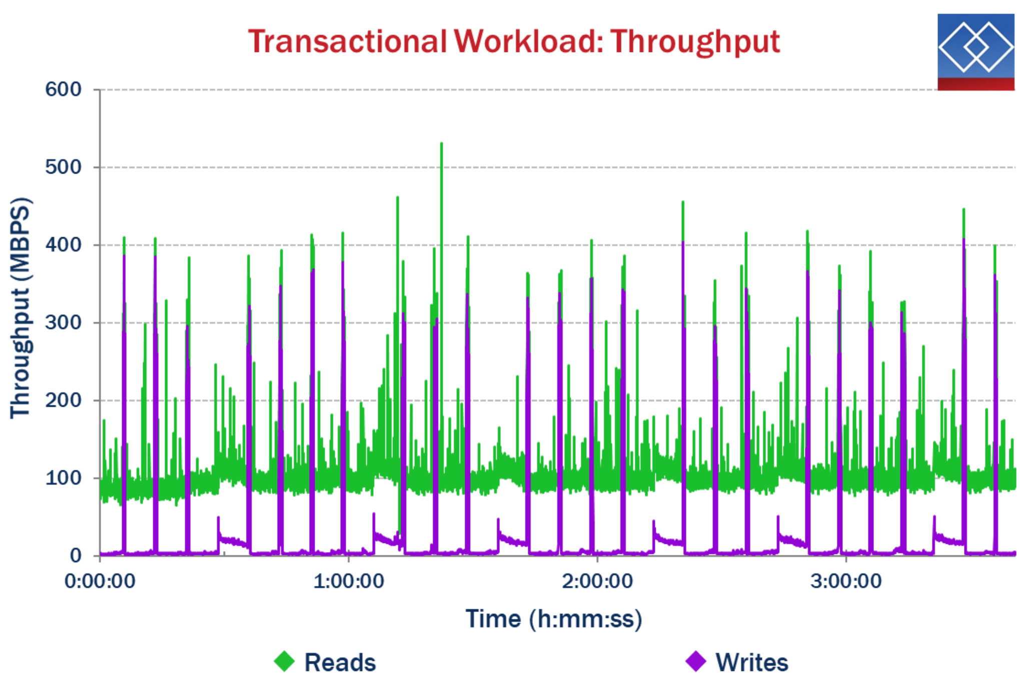

Latency

Although there were periodic spikes due to the workload behavior, the minimum latencies for this environment were quite good.

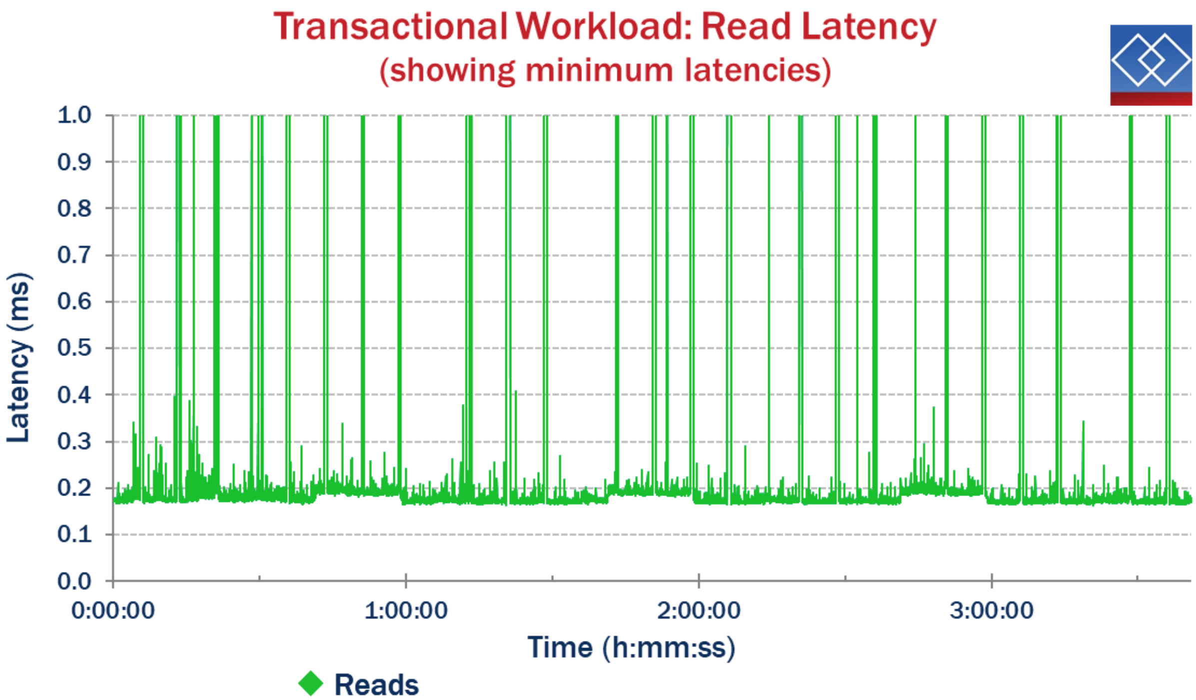

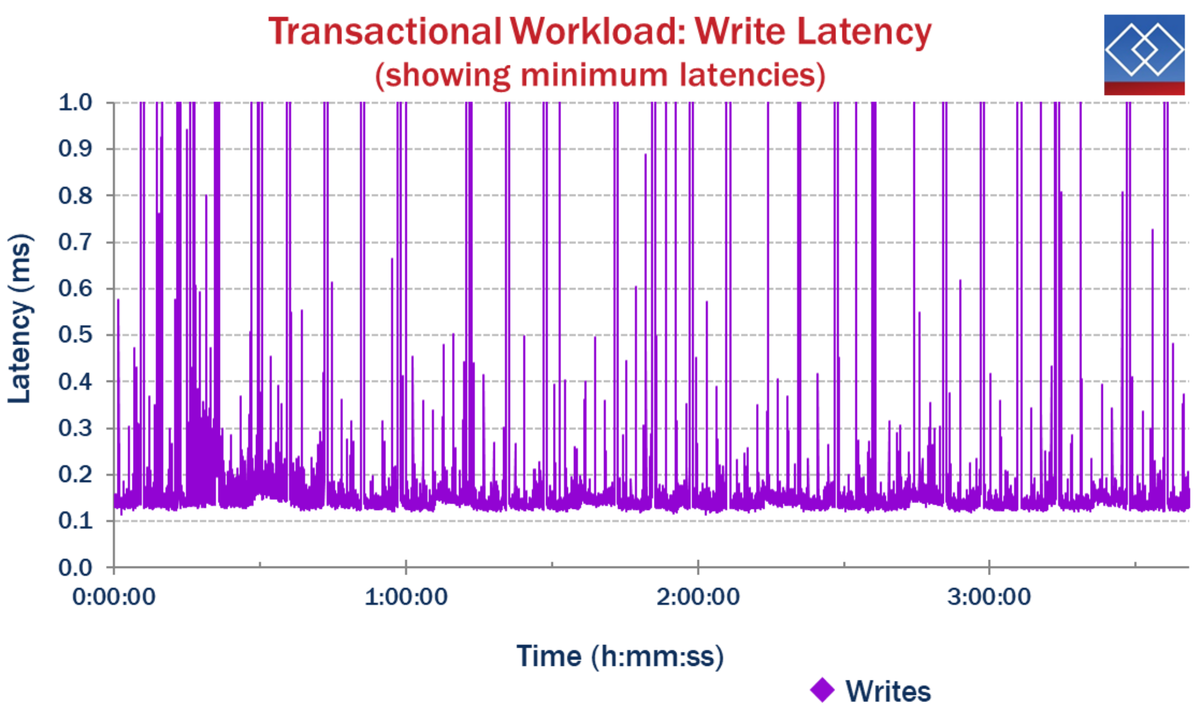

Memory Utilization

Within the OLTP VM, the memory utilization steady.

CPU Utilization

The CPU utilization for this OLTP workload is what we expect, reasonably steady, and not uncommon for this workload.

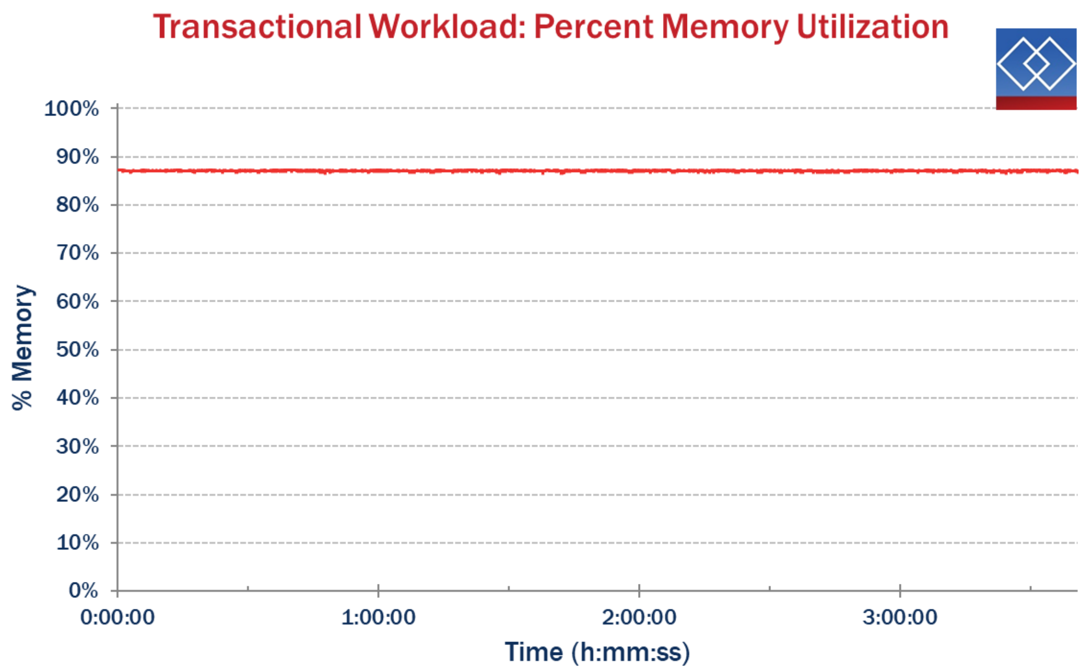

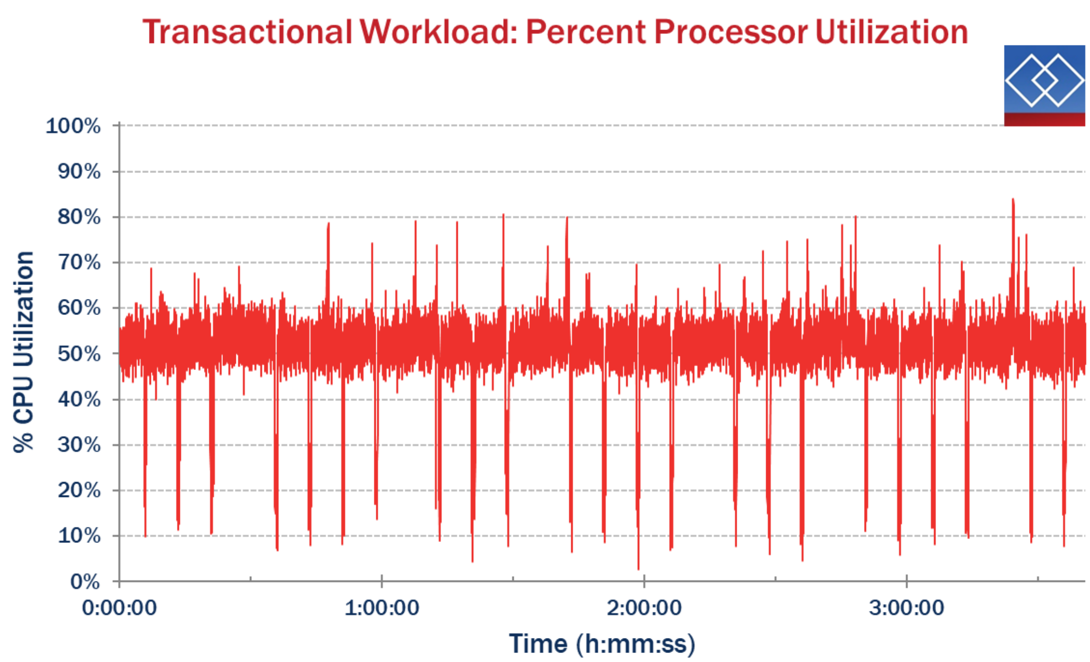

Summary and Conclusion

SAS is a robust technology that has adapted to the changing requirements of data centers for several years, taking advantage of new speed technologies and increasingly advanced features such as dual-ported drives, multi-path I/O (MPIO) support, hot swap ("surprise add/removal"), and more.

In these tests, SAS technology provided very strong performance in a complex, real-world environment.

SAS is a proven and sustainable storage platform that supports current and future storage needs. By leveraging its strong storage legacy, SAS continues to innovate with a roadmap that extends to 24G SAS and beyond.

The most current version of this report is available at http://www.demartek.com/Demartek_SAS_Performance_2018-04.html on the Demartek website.

