Demartek

Dell EMC FX2 Enterprise Database Workloads with Toshiba 12Gb/s SAS SSDs

Tackle Demanding Workloads with Dell EMC FX2, Toshiba PX05S Series 12Gb/s SAS SSDs and Oracle Database

Executive Summary

The Dell EMC PowerEdge FX2 enclosure is a 2U hybrid rack mount computing platform that combines the flexibility, density, and efficiencies of blade server technology with the simplicity and cost benefits of rack-based systems. The FX architecture's innovative modular design supports IT resource building blocks of varying sizes so data centers have greater flexibility to construct their infrastructures. The enclosure uses Toshiba's 12Gb/s SAS SSDs to provide performance and density to the solution. This platform can be attractive to those needing more computing power in a smaller space without having to go through the "learning curve" of blade environments.

We wanted to show a real-world example of the density and simplicity of the FX architecture by deploying an Oracle database solution. For this particular set of tests, we chose two compute nodes and two storage nodes. For the compute nodes, we chose the half-width PowerEdge FC630, a two-socket workhorse server that provides a powerful punch in a small footprint. For the storage nodes, we chose the half-width PowerEdge FD332. These dense storage nodes support 16 small form factor devices – and we filled them with 3.84TB Toshiba PX05S Series 12Gb/s SAS SSDs, providing more than 60 terabytes of flash storage in each storage node. We then assigned one storage node to each compute node.

We deployed the latest version of the Oracle database on both compute nodes, each having its own all-flash storage node. We ran an OLTP workload on both compute nodes to see what kind of performance we could get in a 2U rack space.

Key Findings

> With Toshiba 12Gb/s SAS SSDs, we achieved almost 30,000 database transactions per second on each compute node, for a total of almost 60,000 database transactions per second in one chassis.

> Striping the Oracle database across the Toshiba 12Gb/s SAS SSDs resulted in average latencies of less than 350 microseconds (350 µs).

Dell EMC FX2 Enterprise Database Workloads with Toshiba 12Gb/s SAS SSDs

Demartek
September 2017

Dell EMC PowerEdge FX2 Enclosure

Storage for the FX2 enclosure comes in a 1U half-width form factor. Servers are available in 1U quarter-width, half-width, or full width form factors, depending on processing needs. Application owners can mix and match server and storage nodes to customize a solution that best meets their unique scaling, performance and SLA requirements.

In our example, two Dell EMC PowerEdge FC630 servers and two Dell EMC PowerEdge FD332 storage blocks were included in the FX2 chassis. In two rack units twice the compute and storage power is available compared to a traditional standard 2U rack server. The servers and storage blocks share power, cooling, and management. Each FX2 chassis includes 1100W, 1600W or 2000W power supply units in 1+1 redundant or 2+0 non-redundant configurations. The entire chassis is managed with the same Dell iDRAC interface Dell customers are already familiar with.

Dell EMC PowerEdge FC630 Server

The PowerEdge FC630 packs a powerful processor and ample memory into a compute node that uses only a ½-width unit of rack space. The FC630 is designed to accelerate performance of large databases, virtualization, or mid to large business applications. Each compute node includes:

- 2x Intel® Xeon® E5-2698 v4 processors, 2.2 GHz, 20 total cores, 40 total threads
- Intel C610 series chipset
- 256GB memory. (Each server can take up to 24 DIMMs/1.5TB of DDR4)
- 2x Toshiba 1.6TB PX05S 12Gb/s SAS SSDs for boot drives
- 2 mezzanine PCIe 3.0 (x8) expansion slots for DAS (Slots can also be used for SAN or NAS)

Dell EMC PowerEdge FD332 Storage Block

The FD332 has the flexibility for multiple FD332 to be attached to a single server, for a dual-controller server to split access into two groups of 8 drives, or for pooled storage across multiple controllers and servers.

In our example we used two FD332 Storage Blocks. One FD332 storage block was assigned to each FC630 server. The FD332 uses Direct Attach Storage (DAS), which has low latency as there is no SAN or virtualization layer needed. Each storage block included:

- 16x Toshiba 3.84 TB PX05S 12Gb/s SAS SSDs

One FD332 supplies more than 60 TB per ½-width unit. The two FD332s combined provide 32 drive bays, which is more than is found in most traditional 2U server chassis.

Toshiba PX Series Enterprise SSDs

Both the Dell PowerEdge FC630 Server and Dell PowerEdge FD332 Storage Block used Toshiba PX05S SSDs for optimal performance. The Toshiba SSDs provided more than 100TB in 1U using 2 FD332 Storage Blocks. The Toshiba drives feature:

- 12Gb/s SAS
- (MLC) Multi-Level-Cell NAND flash which reduces cost per bit
- Availability in 1, 3, 10, or 25 (DWPD) Drive Writes Per Day for workload-specific endurance.
- 270K IOPS random read (4K) performance
- Availability in read intensive, mixed use, or write intensive models providing workload specific performance.

The large capacity of the Toshiba 3.84TB PX05S 12GB/s SAS SSDs allowed us to take advantage of RAID 10 for data protection with the smallest write penalty and still have space to spare.

Dell EMC FX2 Enterprise Database Workloads with Toshiba 12Gb/s SAS SSDs

Demartek
September 2017

Oracle Database 12c

Oracle Database 12c delivers performance, scalability, high availability, data optimization, data security and ease of management to support the most demanding OLTP, Data Warehousing and Big Data requirements. Oracle Database 12c offers a variety of features including consolidating databases and enabling Database as a Service with Oracle Multitenant; real time analytics with database in-memory optimizations; automatically optimizing data storage and compression according to usage patterns; providing continuous access with Oracle's Maximum Availability Architecture; securing enterprise data with a comprehensive defense-in-depth strategy, and simplifying the analysis and integration of Big Data and efficient database management using Enterprise Manager Cloud Control.

Oracle Linux

Oracle Linux 7.3 is Oracle's Linux operating system featuring the Unbreakable Enterprise Kernel (UEK) designed for 64 bit Oracle Engineered Systems and Oracle Cloud. Oracle Linux is binary compatible with Red Hat Enterprise Linux (RHEL); all system libraries are identical allowing applications designed for RHEL to run on Oracle Linux.

Test Configuration

Each database was allocated 200GB of system memory, for a total of 400GB of chassis memory dedicated to Oracle.

Database transaction numbers were obtained from Oracle AWR as summarized by HammerDB, while the rest of the data was gathered from each server using the iostat package at 10 second intervals.

HammerDB and OLTP

The HammerDB open source TPC-C implementation was used to generate our OLTP workload. HammerDB runs on a real database, creates real tables, rows, indexes and stored procedures, and executes real queries and transactions, all in the same manner as a bona fide application. The data populating these tables was randomly generated within guidelines defining structure and quantities. Our OLTP workload was modelled on an online retailer taking and delivering orders. Roughly 30% of this workload's transactions are write transactions. The workload was run with a 10 minute ramp-up time and then a 60 minute execution.

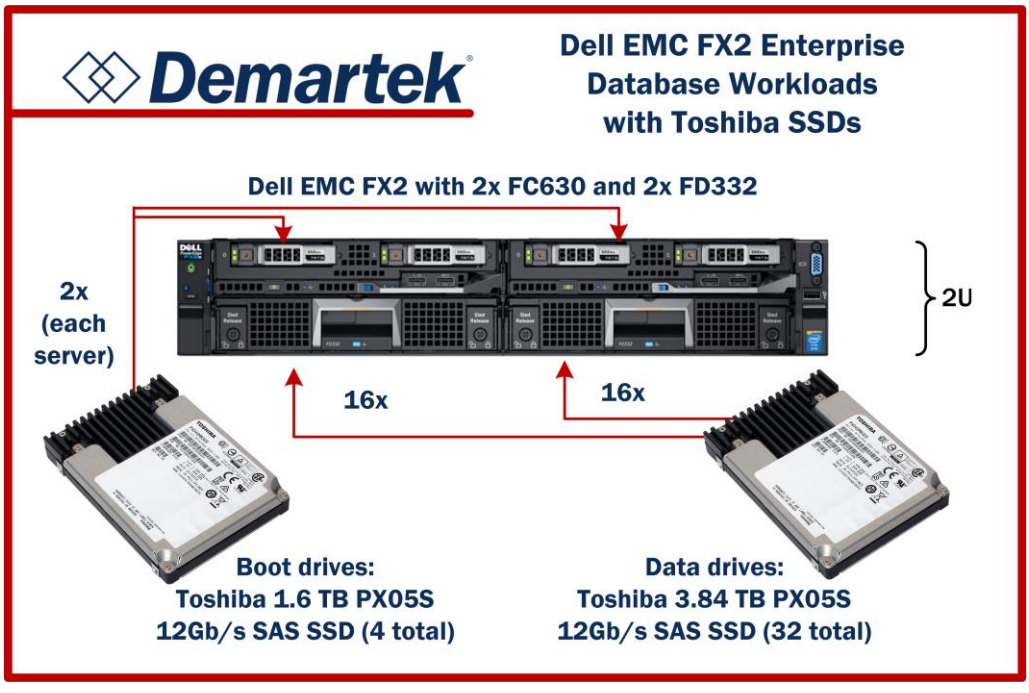

Dell EMC FX2 Enterprise Database Workloads with Toshiba 12Gb/s SAS SSDs

Performance Results

CPU Utilization

CPU utilization is a measure of the amount of CPU processing that occurs when a workload runs on a system. Each system was running a separate OLTP workload, and the average CPU utilization of the servers while running this workload was 53%. The servers were not overly taxed while performing this workload, leaving almost half of their processing power available for other applications or future growth.

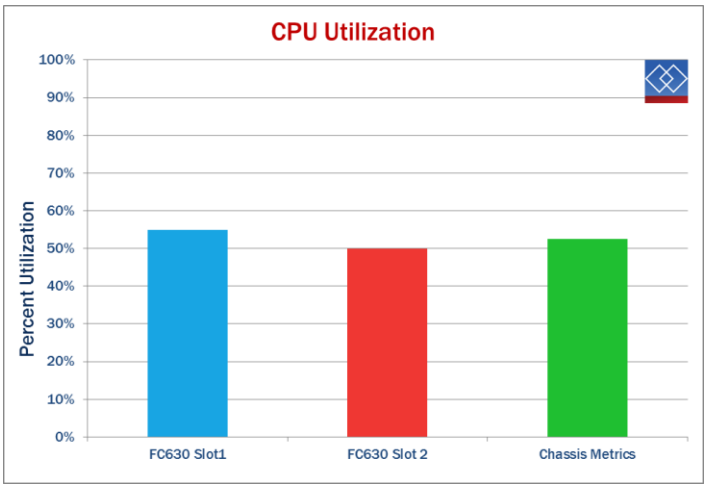

Database Transactions per Second

We achieved nearly 30,000 database transactions per second in each compute node, for a total of nearly 60,000 database transactions in the chassis.

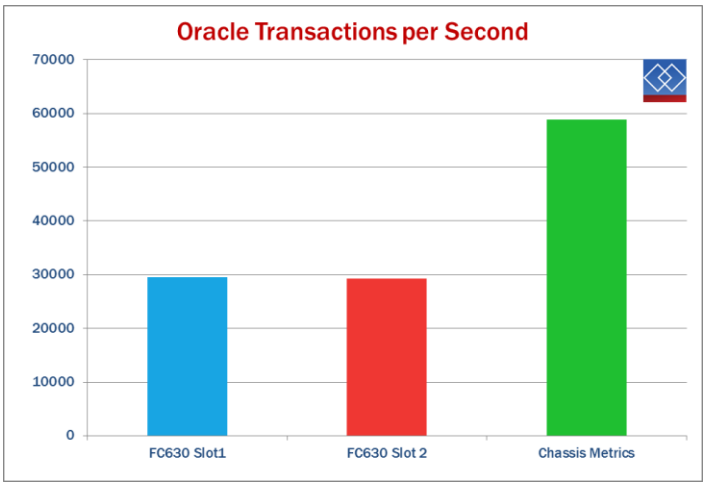

Read and Write Latency

Latency is a measure of the time taken to complete an I/O request, also known as response time. Latency has frequently been measured in milliseconds (ms), which is thousandths of a second. However, with the advent of SSD storage measuring latency in microseconds (µs), which is millionths of a second, is now becoming more common.

Latency can be critical even though it doesn't necessarily have a direct effect on IOPS and throughput. It does have a very significant effect on application performance and user experience. Unlike IOPS and throughput, where more is better, with latency, the goal is to keep it as low as possible.

With the use of an OLTP workload and SSDs, we see the systems maintaining an average of less than 350 microseconds of latency for both reads and writes.

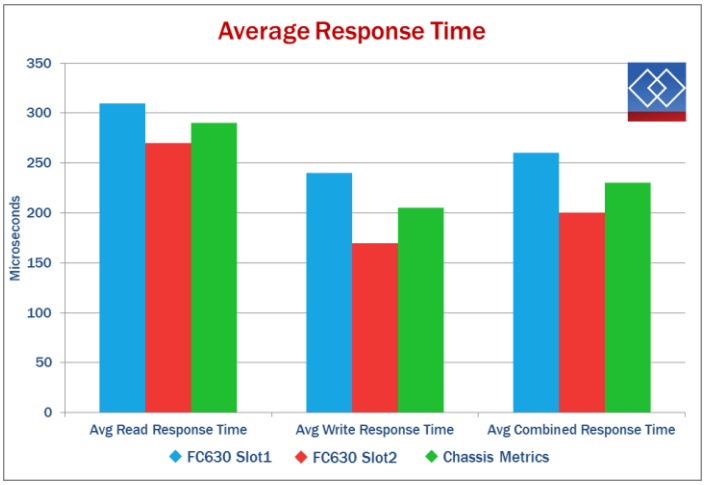

Dell EMC FX2 Enterprise Database Workloads with Toshiba 12Gb/s SAS SSDs

Demartek
September 2017

Summary and Conclusion

Toshiba has been a long-time supplier for Dell, and their partnership continues to benefit customers with performance intensive applications. The Dell FX2 design, with Toshiba 12Gb/s SAS SSDs, would appeal to those who need to get more compute and storage resources into a small amount of rack space, especially those for whom traditional rack servers consume too much rack space.

In 2U of rack space, we had two dual-socket compute nodes and two 16-drive all-flash storage nodes — a very dense package, and were able to achieve almost 60,000 database transactions with an average I/O latency less than 350 microseconds for both writes and reads to the Toshiba PX05S storage media.

The most current version of this report is available at www.demartek.com/Dell-FX2-Toshiba on the Demartek website.