

Technical Validation

Dell EMC PowerMax and SCM Powered by Dual-port Intel Optane Technology Combine to Improve Overall System Performance

Higher Performance and Lower Latency with End-to-end NVMe and Storage Class Memory as Persistent Storage

By Kerry Dolan, Senior IT Validation Analyst

September 2019

This ESG Technical Validation was commissioned by Dell EMC and is distributed under license from ESG.

Contents

Introduction.....	3
Background.....	3
Dell EMC PowerMax	4
Architecture Matters	4
The Difference with NVMe Drives.....	4
The Latest: FC-NVMe Front End and a Persistent SCM Tier Powered by Intel Optane Technology	4
ESG Technical Validation	6
I/O Performance Testing.....	6
Bandwidth.....	6
NAND versus SCM Drive Performance:	7
Application-level Testing.....	7
Mixed Workload Scale Testing.....	8
Mixed Workload QoS Testing	10
The Bigger Truth	12

ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team’s expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

Introduction

This report documents ESG’s audit of performance testing of Dell EMC’s latest PowerMax with end-to-end NVMe and SCM powered by dual-port Intel Optane technology to deliver the high performance and low latency that today’s real-time applications need.

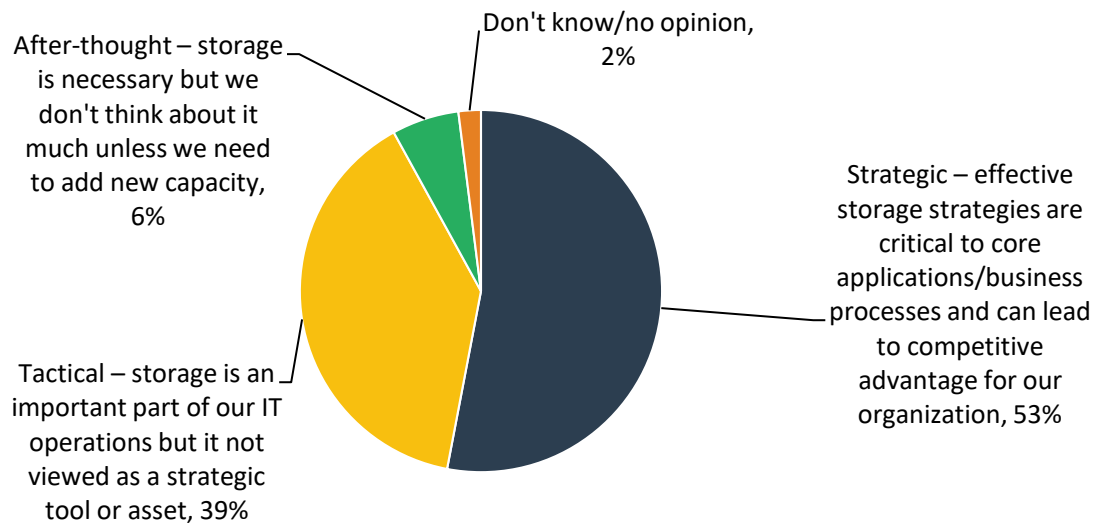
Background

Enterprises are collecting growing quantities of data that are not only used in business-driven applications such as online transaction proceeding (OLTP), Oracle, and SQL Server, but also for customer insight and business trends. In ESG research, 48% of IT decision makers reported both primary production data and secondary data growing by more than 30% annually.¹ Growing business demands and IT complexity make it challenging for IT to design optimized data centers that deliver the levels of performance, scalability, efficiency, and agility required to succeed.

Storage, once considered to be a static, background IT process, today plays a critical role in delivering on business success. The right storage infrastructure can deliver the optimal mix of performance, efficiency, cloud functionality, and cost. When asked what best describes the role that data storage technology plays in their organization’s IT and business operations, 53% of ESG survey respondents agreed that storage was strategic, and that effective storage strategies were critical to their core applications and business processes.²

Figure 1. The Strategic Role of Storage

Which of the following best describes the role that data storage technology plays in your organization's IT and business operations? (Percent of respondents, N=356)



Source: Enterprise Strategy Group

Enterprises need consolidated storage to handle mixed workloads efficiently and cost-effectively, but with the scalability for exponential growth and the high performance/low latency that today’s applications require.

¹ Source: ESG Master Survey Results, [Copy Data Management Trends](#), March 2018.

² Source: ESG Master Survey Results, [2017 General Storage Trends](#), November 2017.

Dell EMC PowerMax

PowerMax is designed for consolidated, mixed workloads that demand high performance at scale, such as OLTP, decision support, real-time analytics, SAP, Oracle, and SQL. The latest PowerMaxOS release brings end-to-end NVMe and SCM drives powered by dual-port Intel Optane technology for optimal performance and low latency; Dell EMC reports exceptional performance of up to 15M IOPS and 350GB/s bandwidth. It is designed for six nines availability, massive scalability (from 13TB up to 4PB effective capacity, 256 front-end ports, 64,000 LUNs), and consolidation of open systems, mainframe, IBM i, file storage, and containerized applications on the same array. Each PowerMax brick includes one engine, up to two array enclosures, and redundant components. PowerMax 8000 supports up to eight bricks (16 controllers), and PowerMax 2000 up to two bricks (four controllers).

Architecture Matters

The latest PowerMax preserves the design concept of EMC's long history of leading enterprise arrays: that is, maintaining a large cache so that hosts access data primarily from cache (DRAM) to deliver the best application performance, but also providing maximum back-end drive speed. PowerMax is built on a multi-controller architecture with a globally consolidated cache accessible to all 16 controllers, enabling most data to be accessed from the cache DIMMs, which provide the lowest latency and speed of absorbing writes. The addition of end-to-end NVMe and Storage Class Memory (SCM) technologies are the latest innovations to support this design.

The Difference with NVMe Drives

While flash drives made a significant impact in storage performance compared with HDDs, those that connect through SATA/SAS interfaces are hindered in I/O parallelization and ultimately in the performance they can deliver. These interfaces were designed for traditional hard drives, not for high speed flash media. The NVMe protocol was developed specifically for non-volatile, high-speed flash media. Key features that enable maximum performance and lowest latency are NVMe's ability to run vastly greater numbers of parallel I/O operations, its streamlined connection to host CPU, and the simpler software stack that reduces I/O processing time and CPU overhead. These features accelerate existing applications and enable new applications that demand real-time processing. In addition, by handling hefty workloads in a smaller footprint, NVMe arrays can reduce TCO.

The Latest: FC-NVMe Front End and a Persistent SCM Tier Powered by Intel Optane Technology

The latest version of PowerMax adds two significant advancements: 1) NVMe over Fabric (NVMe-oF) front-end connectivity that enables end-to-end NVMe; and 2) the option for SCM powered by dual-port Intel Optane as a persistent storage tier for ultra-low latency. These new features provide the storage foundation for mission-critical applications that demand the highest levels of performance, such as OLTP and real-time analytics.

- Adding a 32Gb front-end FC-NVMe I/O module enables end-to-end NVMe for the fastest cache read performance and lowest latency. Enabling the NVMe protocol to run over a Fibre Channel SAN using Connectrix switches and PowerPath makes it easier for organizations to adopt these improvements.
- SCM drives offer performance similar to DRAM volatile memory, but as a persistent media like NAND flash. PowerMax offers the option of adding SCM drives with dual-port Intel Optane SSDs as a persistent storage tier. This not only reduces drive read latency, but also maintains low read response time under the strain of heavy writes to the drives—an improvement over NAND flash, whose read latency suffers with heavy writes. Most organizations will target these drives for specific use cases such as real-time analytics and high-demand OLTP for finance, telecom, retail, fraud detection, etc.

These advancements support the PowerMax architecture concept of making the most use of cache to speed host reads and writes, and reducing the latency when data is accessed from the back-end drives. This powerful combination of the fastest protocol and drive technologies with automated tiering create an array that offers not only maximum performance, but also the ability to balance performance and cost. Workloads that need the highest performance can have it, while other workloads are not burdened with excess cost.

Other PowerMax features include:

- A built-in machine learning engine for automated data placement without overhead. This engine analyzes I/O and uses predictive analytics to optimize performance and cost by assigning data to the best media (flash or SCM) and choosing when to bypass compression/deduplication when datasets are very active to avoid performance bottlenecks. Users can prioritize applications and assign service levels which are incorporated into this automation.
- Linear scalability with NVMe over fabric when scaling from 2-8 bricks, with every port able to access every DIMM and back-end drives.
- Extreme efficiency with global, inline deduplication and compression, including guaranteed 3:1 data reduction through the Dell EMC Future Proof Customer Loyalty Program.
- Data-at-rest hardware-based encryption that is FIPS 140-2 validated and tamper-proof audit logs for security.
- Integrated copy data management.
- Non-disruptive data migrations.
- HTML 5-based Unisphere intuitive management interface.
- Data protection options, including SnapVX space-efficient snapshots; RecoverPoint for any-point-in-time recovery; high availability with SRDF metro; and direct backup to Dell EMC Data Domain.
- Other software options, including CloudIQ storage monitoring and analytics; PowerPath multi-pathing; Storage Resource Management metrics and reporting tools.

Figure 2. PowerMax Q3 2019 Release



- ✓ Up to 4PB effective capacity
- ✓ Up to 15M IOPS
- ✓ Up to 350 GB/s throughput
- ✓ Up to 16TB Cache
- ✓ Up to 576 CPU cores
- ✓ Infiniband Virtual Matrix
- ✓ NVMe back end
- ✓ Up to 256 server Ports: FC, FICON, iSCSI, **NVMe-OF**
- ✓ Up to 288 drives: NVMe Flash or **SCM**



Source: Enterprise Strategy Group

ESG Technical Validation

ESG audited the results of testing conducted at Dell EMC’s Hopkinton, MA facilities. Testing was conducted with Dell EMC internal tools to demonstrate overall performance using the new 32Gb FC director and SCM drives; application-level testing with Oracle and SQL Server demonstrated mixed workload scalability and quality of service (QoS) using industry standard benchmarks.

I/O Performance Testing

ESG audited the results of Dell EMC performance testing that measured IOPS, latency, and bandwidth using I/O generated from Dell EMC PowerEdge R740 servers running SLES and Windows Server 2012. An overview of our audit is provided below.

Bandwidth:

Port bandwidth testing used a two-engine PowerMax 8000 and compared the performance of 16Gb and 32Gb front-end directors using a 128K I/O workload of random read and random write cache hits. Single- and dual-port testing demonstrated linear scalability for reads and near-linear for writes (additional ports ran into a PCI limit).

- For random reads, a single port scaled 100% from 1.53 Gbps to 3.07 Gbps; for random writes, a single port scaled from 1.44 Gbps to 2.74 Gbps.
- For random reads, dual ports scaled 100% from 3.07 Gbps to 6.14 Gbps, and for random writes, dual ports scaled from 2.87 Gbps to 5.49 Gbps.

PowerMaxOS Q3 2019 Release With End-to-end NVMe and SCM Drives



- ✓ 500% more write IOPS
- ✓ 100% more bandwidth
- ✓ 26% lower latency

NAND versus SCM Drive Performance:

Back-end drive testing used a single-engine PowerMax 8000 and compared NAND flash with SCM drive performance. Workloads were configured to miss cache to ensure back-end drives were tested.

- For an 8K random read I/O workload using NVMe-NAND drives, response time at 100K IOPS was 0.2921 ms. Using NVMe-SCM drives with Intel Optane SSDs, response time at 100K IOPS was 0.2162 ms, 26% lower latency.
- A write intensive test using 8K random reads and 128K random writes demonstrated that NAND flash drives delivered up to 5K IOPS, while SCM with Intel Optane SSDs delivered up to 30K IOPS, 500% more than NAND IOPS.



Why This Matters

To remain competitive, enterprise IT organizations today must deliver fast performance to support the needs of applications in real time.

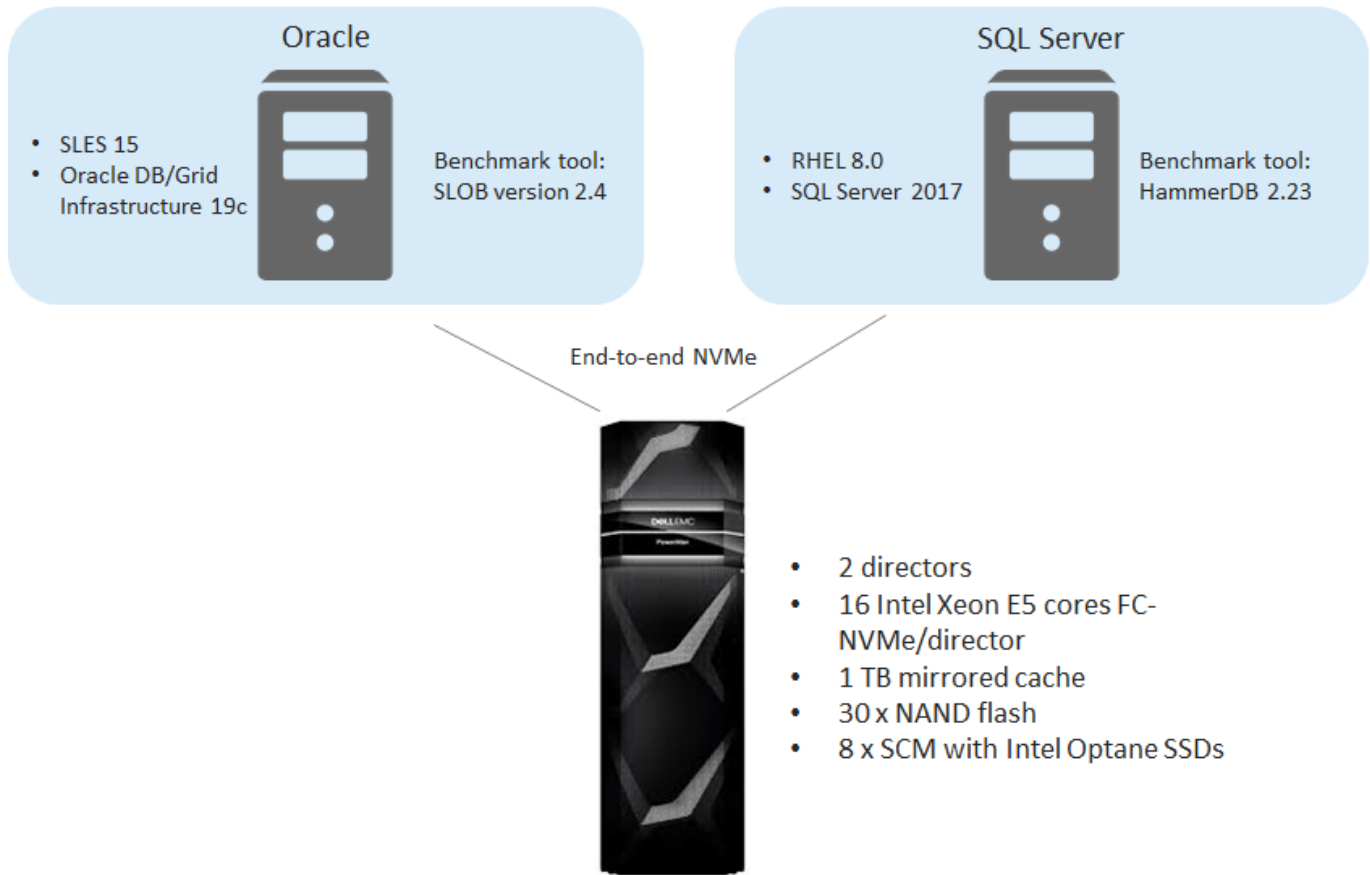
ESG validated that the latest single-engine PowerMax 8000 with a 32Gb FC-NVMe director and SCM with dual-port Intel Optane drives provided 500% more write IOPS, 100% more bandwidth, and 26% lower latency than the 16Gb director and NAND flash drives. These capabilities enable organizations to get more work done in the same footprint, reducing TCO with consolidated management and lower node-based licensing costs for applications like Oracle.

Application-level Testing

ESG validated application-level testing of both scale and QoS. The test bed used a single engine PowerMax 8000 with two 32Gb FC-NVMe directors, 1TB raw mirrored cache, and storage capacity provided by 30 NAND flash drives and eight SCM drives powered by dual-port Intel Optane technology. Hosts for these tests were two 24-core Dell EMC PowerEdge R740 servers, each with 256 GB of RAM, running PowerPath 7.0

- Oracle workloads were executed with the publicly available Silly Little Oracle Benchmark ([SLOB](#)) release 2.4. SLOB is an Oracle I/O workload generation tool kit that is used to measure the suitability of a hardware platform for Oracle database deployments requiring high performance I/O. SLOB tests were run on SLES 15 and Oracle database/Grid Infrastructure version 19c.
- SQL Server workloads were executed using [HammerDB](#) version 2.23, an industry-standard, open-source database load testing and benchmarking tool. The OLTP workload tested with HammerDB emulates the activity of users in a typical online brokerage firm as they generate trades, perform account inquiries, and execute market research. HammerDB tests were run on RHEL 8.0 and SQL Server 2017.

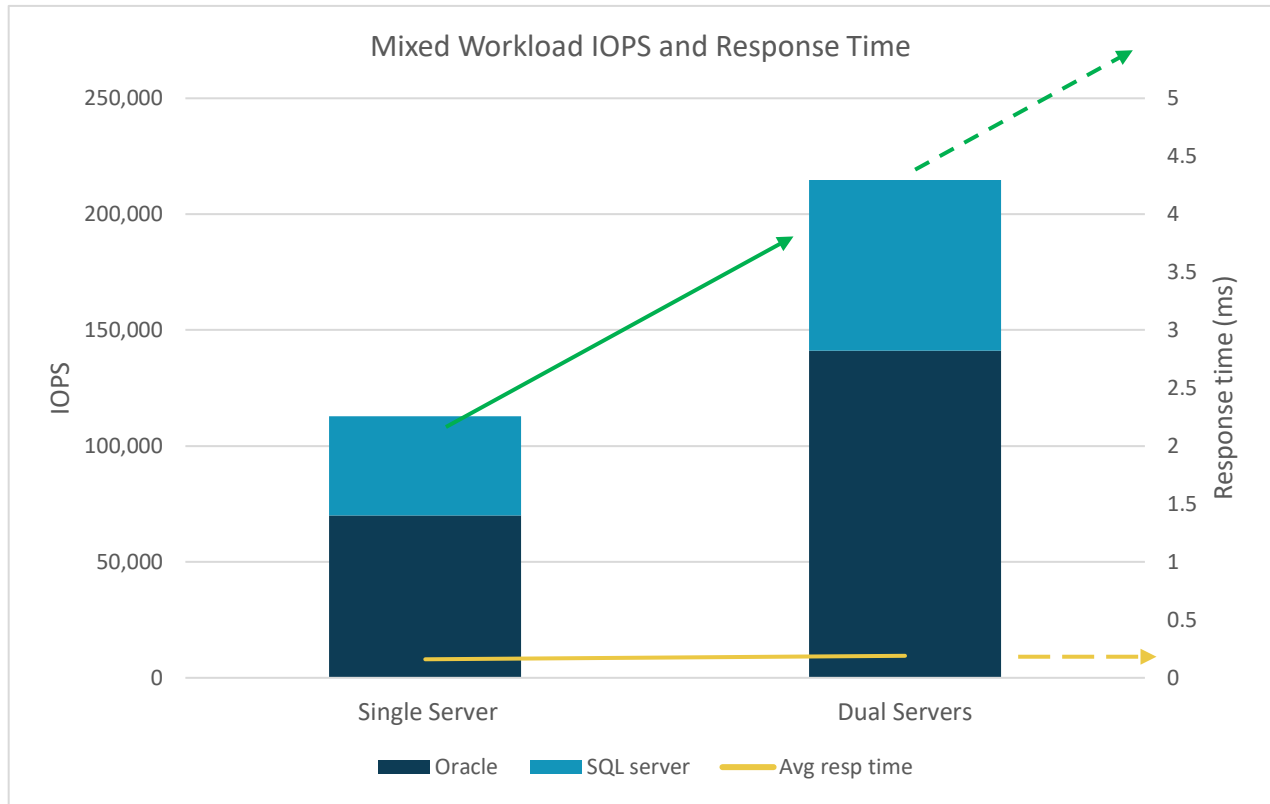
Figure 3. Oracle and SQL Server Test Bed



Source: Enterprise Strategy Group

Mixed Workload Scale Testing

First, we reviewed the mixed workload scale test results. Oracle and SQL Server workloads were generated first with a single server and then two servers with conservative cache hits of 60% for Oracle and 45% for SQL Server. Single server testing demonstrated sub-millisecond response times of 0.19 ms and 0.13 ms, providing extremely low application latency. With two servers, IOPS increased 101% for the Oracle workload and 72% for SQL, but response time remained extremely low at 0.19 ms and 0.23 ms. Figure 4 shows the IOPS and response times of each workload during scaling. The dotted lines indicate that latency would remain low even with an additional increase in IOPS if PowerMax bricks and servers were added.

Figure 4. Scale Test


Source: Enterprise Strategy Group

Table 1. PowerMax Scale Test Details and Application-level Metrics

Test	Workload	IOPS	I/O read response time (ms)	Log file parallel write Avg wait (ms)	TPM
Single Server	Oracle	70,128	0.19	0.37	N/A
	SQL Server	42,604	0.13	N/A	339,018
Dual Servers	Oracle	141,176	0.14	0.38	N/A
	SQL Server	73,485	0.23	N/A	602,343

Additional Application-level Metrics

To make these results more applicable in the real world, we reviewed specific Oracle and SQL metrics. The log file parallel write metric is an Oracle Automatic Workload (AWR) statistic that quantifies the average wait time for processes that are waiting for blocks to be written to the online redo log. Significant log file parallel write wait times are a likely indicator of an I/O issue with the Oracle Log file. As Table 1 shows, log file write times of under half a millisecond remained essentially constant as the amount of database activity doubled during testing (0.37 ms to 0.38 ms).

- This Oracle application-level metric, combined with the high level of IOPS and low read and write response times, is a clear indicator that there were no noticeable I/O bottlenecks as the amount of database traffic doubled.

TPM is a measure of the new-order transactions per minute that were completed during the SQL Server OLTP test. Transactions per minute is a good indicator of the end-to-end performance capabilities of a SQL Server database and its underlying hardware infrastructure. In this case, the number of transactions per minute rose by 78% when the amount of OLTP traffic doubled.

- This correlates well with the 72% IOPS increase and is a good indicator of the predictable scalability of the single node PowerMax that was tested.

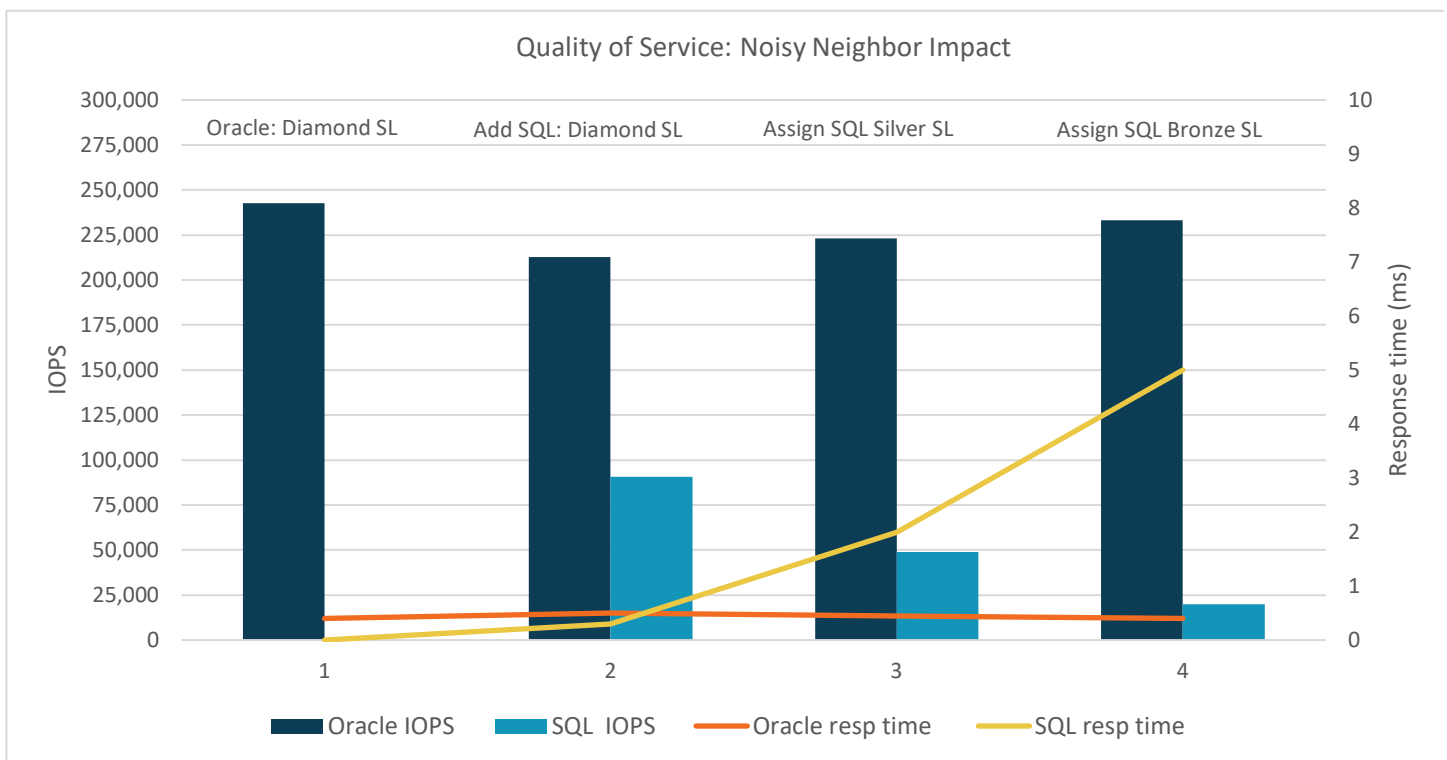
Using the results of this single node performance test as the basis for a theoretical projection of a maximally configured PowerMax platform, ESG is confident that the performance of the PowerMaxOS Q3 2019 release can support a mixed workload that includes business-critical Oracle applications and a virtualized SQL Server infrastructure serving hundreds of thousands of customers and millions of transactions per minute.

Mixed Workload QoS Testing

Next, we reviewed results of QoS testing that demonstrates the performance of applications when “noisy neighbor” applications compete for resources. The test occurred in steps:

- Step 1: Oracle workload assigned a Diamond service level to get the highest priority performance.
- Step 2: Add SQL Server workload with no service level assigned.
- Step 3: Assign SQL Server workload a second tier Silver service level.
- Step 4: Assign SQL Server workload a third tier Bronze service level.

Figure 5. QoS Test: IOPS and Response Time



Source: Enterprise Strategy Group

Table 2. PowerMax QoS Test Details

Test Step	Oracle IOPS	Oracle RT (ms)	SQL IOPS	SQL RT (ms)
Step 1	242,720	0.40	N/A	N/A
Step 2	212,633	0.50	90,665	0.3
Step 3	223,004	0.45	48,978	2.0
Step 4	233,171	0.40	20,129	5.0

Figure 5 and Table 2 show that Oracle IOPS began at 242K and dropped 12% when the SQL workload was added, but then increased 5% when the SQL service level was reduced to Silver and another 5% when the SQL service level was reduced to Bronze. Oracle response time remained extremely low (no higher than half a millisecond) throughout the test.

SQL IOPs started at 90K when unrestrained by service level but dropped 46% when assigned a Silver service level and another 59% when that service level dropped to Bronze. SQL response time started low at 0.3 ms but increased to two milliseconds when the service level dropped to Silver and to five milliseconds when it dropped to Bronze.

This testing demonstrates that the high-priority workload retained high IOPS and low latency while the lower priority workload continued functioning with lower performance.



Why This Matters

Enterprise applications like Oracle and SQL Server drive businesses every day and demand high performance for workloads such as OLTP and real-time analytics. To collect the massive data volumes needed and deliver virtually instant data access for applications such as online trading, fraud detection, retail customer interactions, and the like, these applications must be supported by an enterprise-class storage foundation. Large-scale consolidation of mixed workloads can increase cost efficiency as long as mission-critical workloads are assured of the highest performance.

ESG validated the ability of the latest PowerMax 8000 with end-to-end NVMe and SCM powered by dual-port Intel Optane drives to maintain extremely low latency (less than a quarter of a millisecond) for mixed Oracle and SQL Server workloads while IOPS increased. Oracle log files indicated no I/O bottleneck, and SQL transactions-per-minute scaled in-line with increased IOPS. These results, and the field-proven scalability of the PowerMax as engines are added, confirms that the PowerMax can serve hundreds of thousands of customers and millions of transactions per minute.

We also validated the ability of this array to provide priority performance based on service-level assignments in mixed workload environments. Lower priority applications can run simultaneously, but a high-priority workload will retain its IOPS and response time regardless of other activity. This enables consolidation on a single array without compromising application performance.

The Bigger Truth

Infrastructure vendors must continue to innovate in order for enterprises to take full advantage of emerging technologies. It was not long ago that applications such as real-time analytics were considered bleeding edge, only for the most tech-savvy organizations with the deepest pockets, but advancements have made these technologies available to almost any organization. For many, the challenge now is to create an infrastructure that can support the massive data volumes and fast data access that are required for analytics, fraud detection, trading, interactive retail, and future technologies still on the drawing board.

It should come as no surprise that Dell EMC has made great strides to deliver on these requirements, as the company is well known for its investments in R&D. The latest PowerMax OS release has added end-to-end NVMe and SCM powered by dual-port Intel Optane technology to the already fast, scalable storage platform, and, unlike the few other arrays that support SCM, PowerMax includes SCM drives as a persistent storage tier, not as cache. These new capabilities offer significant performance improvements while enabling organizations to consolidate block, file, and mainframe workloads, shrinking the hardware footprint for reduced power consumption, streamlined management, and lower TCO.

And the performance improvements are quite real; Dell EMC reports exceptional performance of up to 15M IOPS and 350 GB/s throughput with under 100 microseconds read latency. ESG validated both internal and application-level testing with a single-engine PowerMax 8000. The first tests demonstrated 500% more write IOPS, 100% more bandwidth, and 26% lower latency. Oracle and SQL Server testing demonstrated extremely low response times—less than a quarter of a millisecond—even as IOPS doubled, along with the ease with which PowerMax prioritized performance for a mission-critical Oracle workload while running a SQL Server workload at various service levels. These are levels of performance with persistent storage that begin to rival the speed of memory.

ESG looks forward to seeing how customers respond to these new PowerMax features. At this point, the dual-port Intel Optane SCM drives have not seen wide industry testing, but Dell EMC has invested heavily in testing and integration to ensure smooth operation. In addition, NVMe-oF is only available on Linux so far, but support for other operating systems will be available in the near future.

PowerMax is the culmination of years of Dell EMC innovation in storage architecture. It is flexible and scalable, with AI/ML-driven data placement and a future-proof design. With its multi-controller, “shared everything” scale-out architecture, end-to-end NVMe, and SCM with Intel Optane SSDs, PowerMax offers the highest levels of performance, scalability, consolidation, and efficiency for both traditional high-demand workloads such as Oracle, SQL Server, and other OLTP and next-generation applications such as real-time analytics.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2019 by The Enterprise Strategy Group, Inc. All Rights Reserved.

