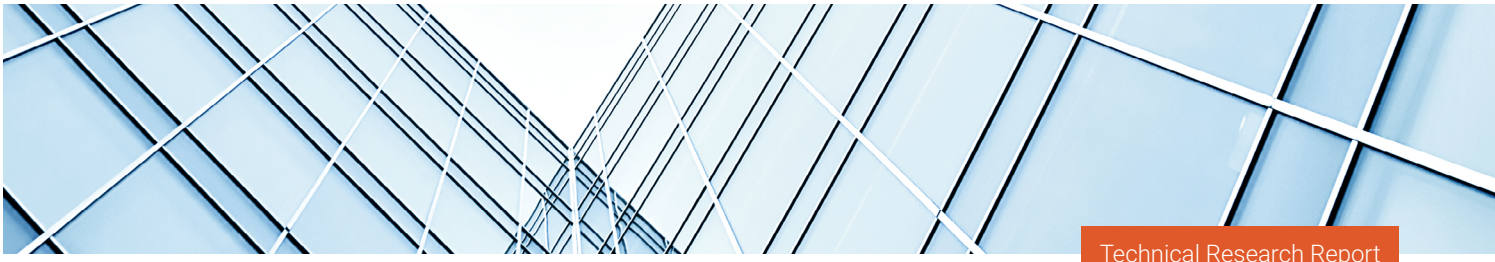

Technical Research Report

PROWESS

Improve Storage Performance by More than 10x for Data-Intensive Workloads on Dell™ PowerEdge™ R760 Servers with Dell™ PowerEdge RAID Controller 12 (PERC 12)

Prowess Consulting testing demonstrated that next-generation PowerEdge servers with the latest NVM Express® (NVMe®) RAID controllers can accelerate business workloads.

Executive Summary

If your business strategy depends on running analytics, AI, or real-time transaction processing at speed, you depend on high performance from your storage. At the same time, it's essential to maintain availability for your storage.

RAID can be a powerful tool for ensuring the availability of local storage by combining multiple physical disk drives into a single virtual disk for the purpose of data redundancy, performance improvement, or a combination of both. However, for many years high-performance storage based on the NVM Express® (NVMe®) protocol lacked native RAID controller support and thus could not directly take advantage of RAID's many benefits. This situation changed with the introduction of newer RAID controllers, such as Dell™ PowerEdge RAID Controller 11 (PERC 11) and PERC 12, which brought native NVMe RAID support, unlocking the ability to directly harness RAID's benefits for high-performance NVMe storage.

To help organizations evaluate the performance benefits of running data-intensive workloads on servers with NVMe RAID, and to assess the benefits of upgrading to Dell PERC 12 from PERC 11, Prowess Consulting tested a Dell™ PowerEdge™ R760 server with both PERC 12 and PERC 11. Our testing found that using PERC 12 instead of PERC 11 allowed us to improve storage performance more than 10x and increase storage throughput almost 5x. Additionally, with PERC 12, RAID 5 rebuilds were up to 110% faster, RAID 10 rebuilds were up to 56% faster, and Microsoft® SQL Server® database restores were up to 90% faster.

As part of our analysis, we evaluated other features of the PowerEdge R760 server that can also contribute to higher performance and that should be considered as part of an upgrade decision. Finally, taking a holistic view of Dell™ servers, PERC, tools, and services, we determined that—for organizations looking to use next-generation Dell servers—the Dell Technologies portfolio can increase server utilization and simplify management, which can reduce both complexity and total cost of ownership (TCO).

Highlights

Data-intensive workloads running on servers with NVMe® RAID deliver higher performance with Dell™ PERC 12.

>>

Disk performance with PERC 12 was more than **10.8x higher** than with PERC 11.

Disk throughput with PERC 12 was more than **4.7x higher** than with PERC 11.

RAID 5 rebuilds on PERC 12 were up to **110% faster** than on PERC 11.

RAID 10 rebuilds on PERC 12 were up to **56% faster** than on PERC 11.

Microsoft® SQL Server® database restores were up to **90% faster** than on PERC 11.

Unleash the Performance of NVMe RAID

RAID enhances the performance, reliability, and data availability of storage systems by combining multiple hard drives into an array. This technology can improve performance by increasing the number of drives used for saving and accessing data. RAID disk subsystems can help improve input/output (I/O) performance and data availability, and they can enhance throughput, as several disks can be accessed simultaneously. RAID systems can also be used to improve storage availability and fault tolerance.

A RAID controller, also known as a RAID card or RAID adapter, is a hardware device or integrated component that manages the RAID array. Although past RAID controllers were not NVMe compatible with the high-speed PCIe® interfaces used by NVMe drives, recent RAID controllers from Dell Technologies (for example) *are* compatible in that respect.

Quantifying the Benefits of a PERC Upgrade

To help organizations evaluate the performance benefits of running data-intensive workloads on servers with NVMe RAID, we tested a PowerEdge R760 server with PERC 11 and the same server with PERC 12. See the **Appendix** for server configurations.

Measuring Disk Performance

Maintaining disk throughput is essential for modern businesses. When disk throughput is constrained, storage can become a bottleneck that can impact the speed of business operations and ultimately reduce profitability.

Prowess Consulting engineers tested the disk throughput of the PowerEdge R760 server platform. For this testing, we used fio, RocksDB, and Dbench (part of the Phoronix Test Suite) with one PowerEdge R760 server using PERC 12 (Broadcom®/LSI PERC 12 H965i Front) and the other one using PERC 11 (Broadcom/LSI PERC 11 H755N Front). Across benchmarks, PERC 12 in the PowerEdge R760 server platform outperformed the same platform using PERC 11. These results can serve as a proxy indicator of how PERC 12 can increase storage system performance in other data-intensive workloads including analytics and AI.

Fio

Fio is a versatile I/O testing tool that benchmarks the performance of storage devices such as the solid-state drives (SSDs) in the PowerEdge R760 server platform at a lower level, independent of any specific application like a database. It tests the storage subsystem's performance by generating various types of I/O workloads, allowing for a detailed analysis of read/write speeds, I/O operations per second (IOPS), and latency under different conditions. Figure 1 and Figure 2 show that PERC 12 delivered more than 10x disk performance for the PowerEdge R760 server platform with both RAID 5 and RAID 10 storage configurations, compared to PERC 11.[2]

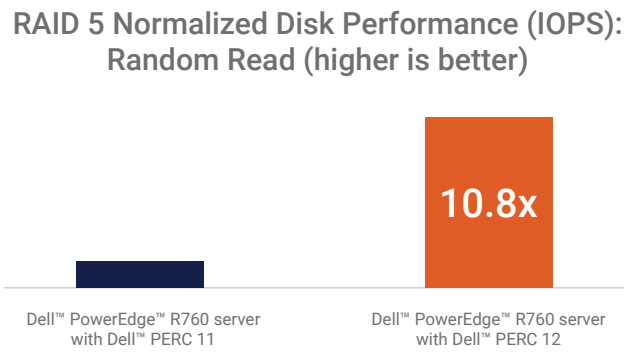

Figure 1 | Dell™ PERC 12 delivers up to 10.8x better RAID 5 fio random read performance

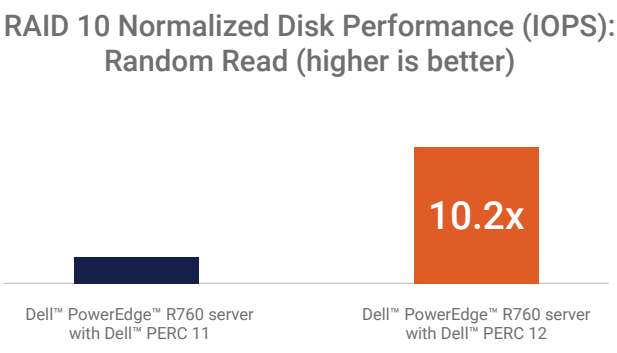

Figure 2 | Dell™ PERC 12 delivers up to 10.2x better RAID 10 fio random read performance

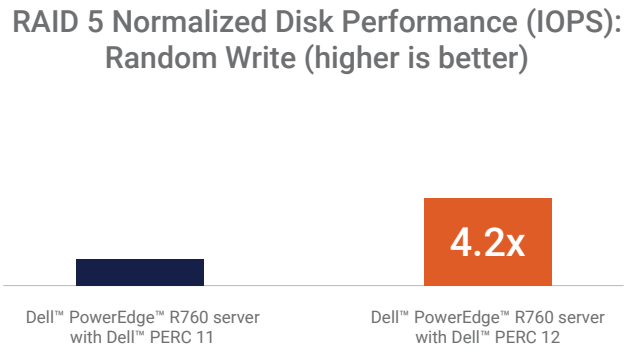

Figure 3 | Dell™ PERC 12 delivers up to 4.2x better RAID 5 fio random write performance

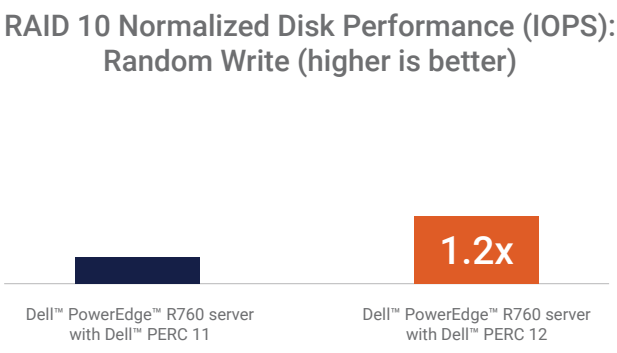

Figure 4 | Dell™ PERC 12 delivers up to 1.2x better RAID 10 fio random write performance

RocksDB

The RocksDB benchmark is specifically designed to measure metrics like throughput, latency, and IOPS within the context of database operations. The benchmarks focus on database-specific operations such as read, write, and compaction processes, reflecting how the database performs under various workloads and configurations.

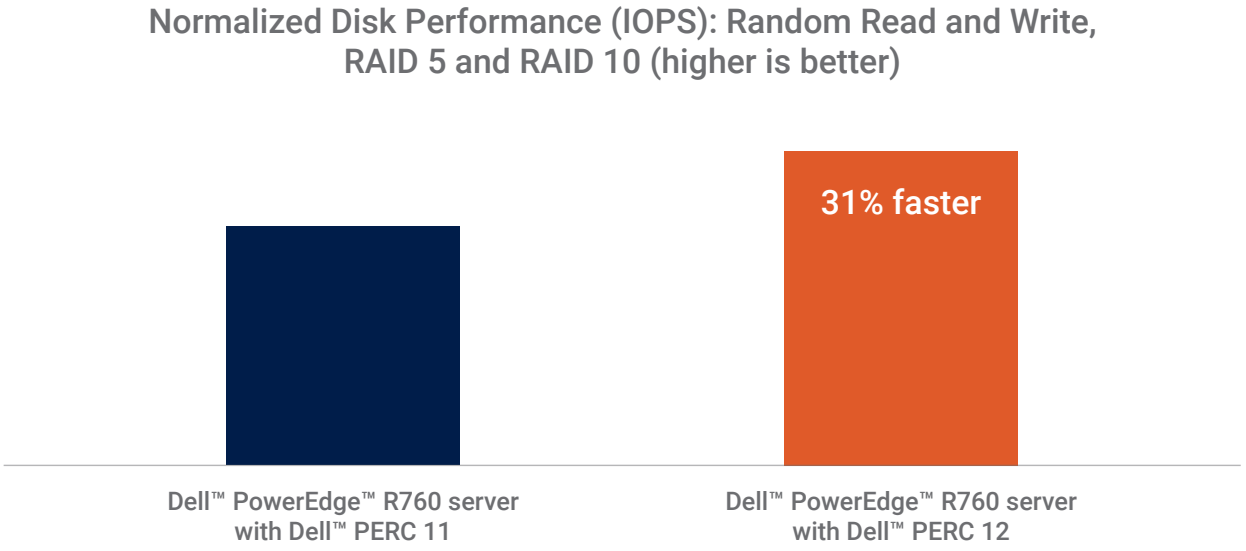

Figure 5 | Dell™ PERC 12 demonstrated higher disk performance with storage configured in both RAID 5 and RAID 10, compared to Dell PERC 11, as measured with the RocksDB benchmark

Dbench

The Dbench benchmark is designed to simulate disk and file system load. It generates a series of file I/O operations, mimicking the behavior of varying numbers of simulated clients accessing a server's file system concurrently (1, 6, 12, 48, 128, and 256 clients across Prowess Consulting's test runs). The PowerEdge R760 server with the newer PERC 12 achieved more than 4.7 times the disk throughput with RAID 5 and more than 1.5 times the disk throughput with RAID 10, compared to the PowerEdge R760 server with PERC 11 for both RAID configurations (see Figure 6 and Figure 7), as measured with Dbench for one client.

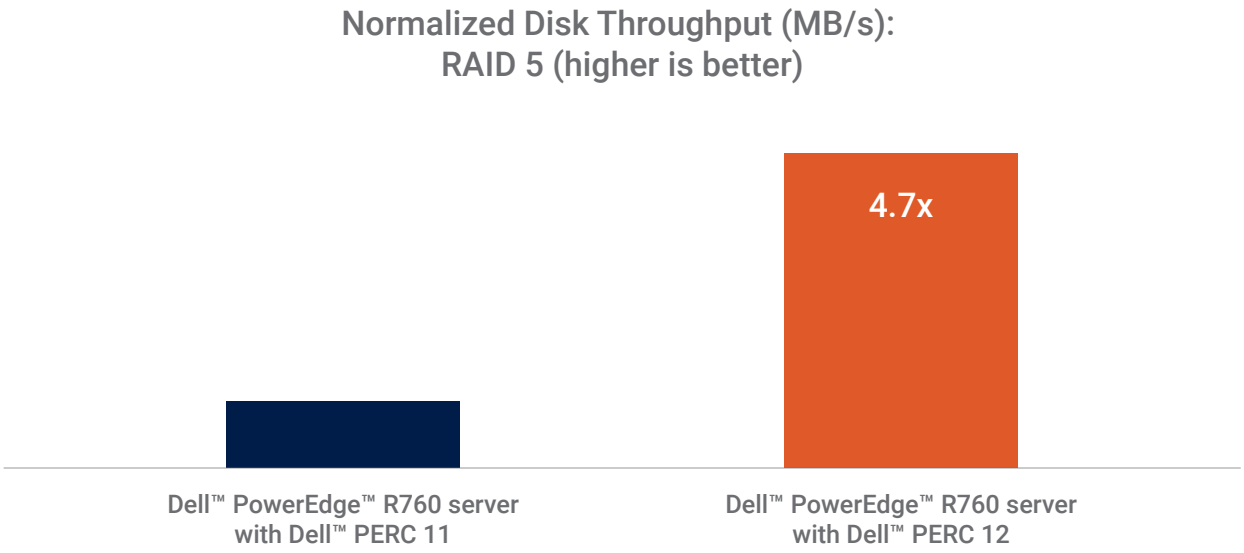

Figure 6 | Dell™ PERC 12 demonstrated higher disk throughput with storage configured in RAID 5, compared to Dell PERC 11, as measured with the Dbench benchmark

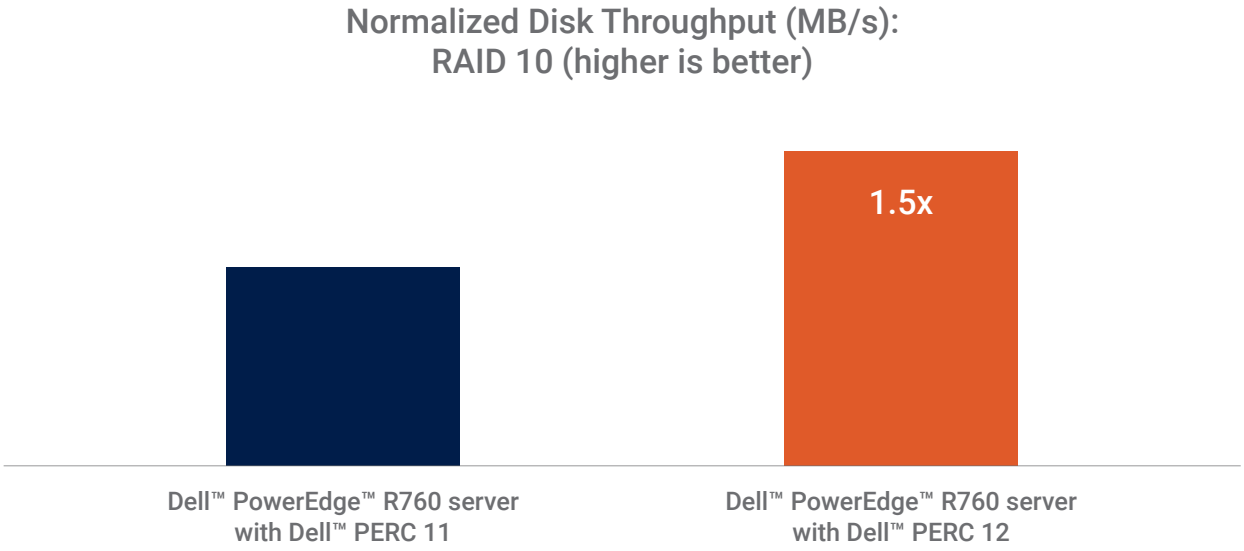

Figure 7 | Dell™ PERC 12 demonstrated higher disk throughput with storage configured in RAID 10, compared to Dell PERC 11, as measured with the Dbench benchmark

Additional RAID Benefits

Using RAID with NVMe drives can also provide other benefits. For example, with SQL Server, the transaction log acts as a buffer for the main database. That means administrators can optimize servers for performance by configuring the front RAID controllers to use RAID 10, which is the highest-performing option. For the back-end database, which is not as performance-sensitive, admins can choose to use RAID 5 (or any preferred level) to sacrifice some performance for increased capacity.

Rebuild Times Drop

When platforms rely on RAID arrays, rebuild times can come into play after a drive failure or replacement; while a RAID array is being rebuilt, there's a risk of permanent data loss if another drive fails before full redundancy has been re-established. As a result, rebuild times can be critically important to organizations. To determine whether the newer RAID array offered an advantage in this area, Prowess Consulting compared the rebuild times of PERC 12 to those of PERC 11. To perform this testing, engineers removed one of the drives from the data array and then recorded the time required to rebuild the system and resume SQL Server production use. Figure 8 shows the median of the two runs for each array. For RAID 5, the PERC 12 rebuild was 110% faster than PERC 11. For RAID 10, the PERC 12 rebuild was 56% faster than PERC 11.

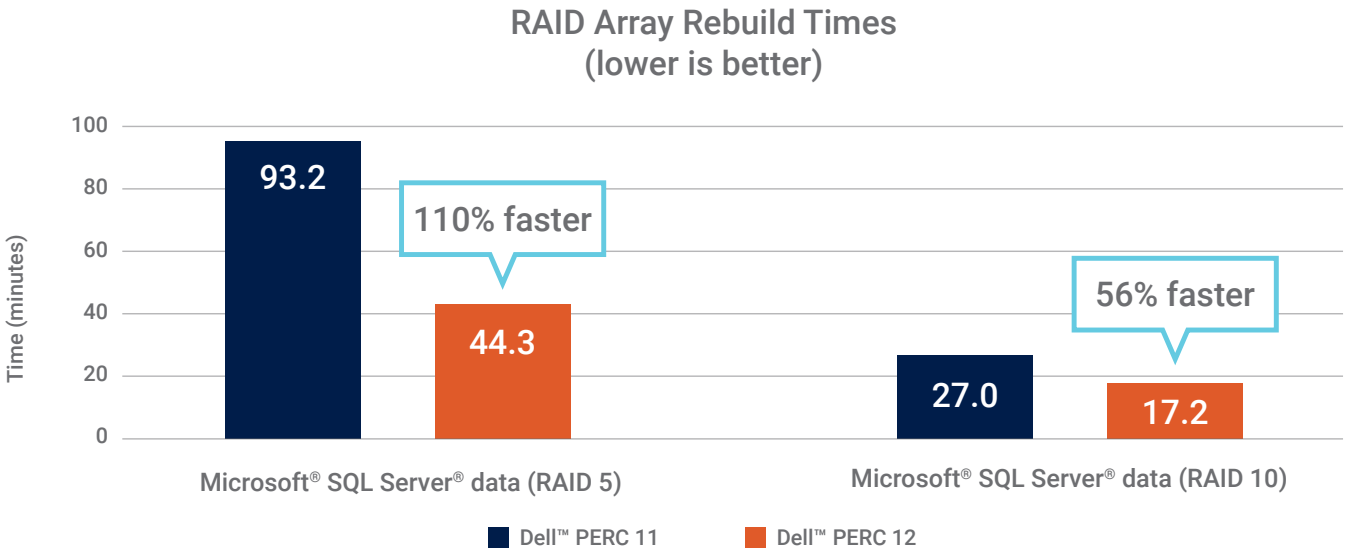

Figure 8 | Dell™ PERC 12 demonstrated faster RAID-array rebuild times for both RAID 5 and RAID 10, compared to Dell PERC 11

Database Restore Times Drop

Like rebuild times, database restore times are important. If a database becomes corrupted, it can be restored from a RAID array, and the quicker, the better. Prowess Consulting compared the Microsoft SQL Server database restore times of PERC 12 to PERC 11. Figure 9 shows that the database restore time for PERC 12 was 90% faster than the restore time for PERC 11.

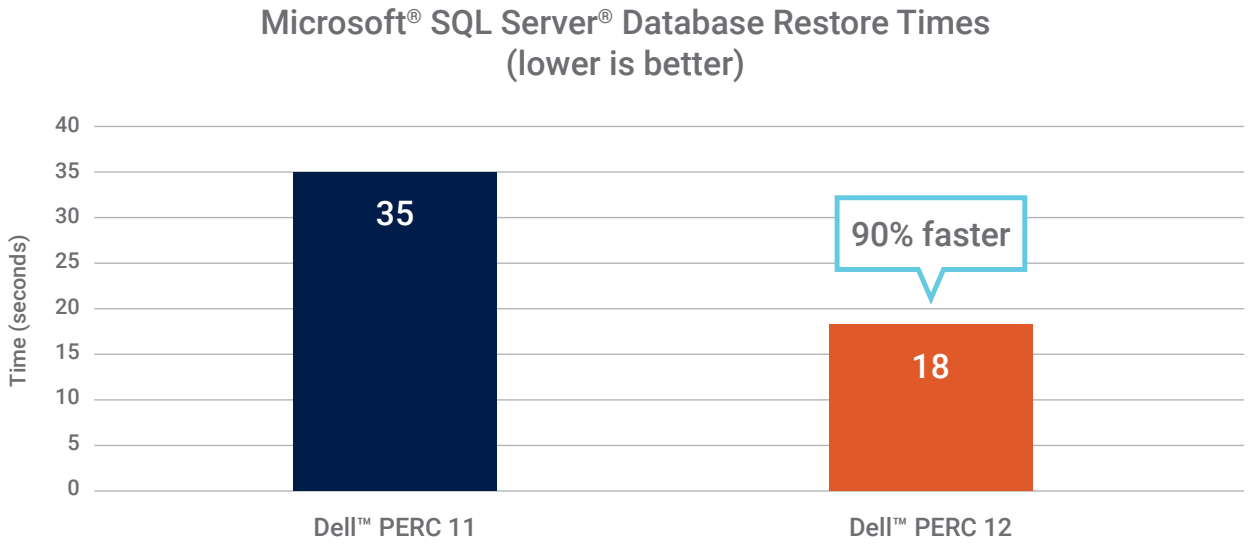

Figure 9 | Dell™ PERC 12 demonstrated faster database restore time for Microsoft® SQL Server®, compared to Dell PERC 11

The Upgrade Decision

Organizations that build their business strategies around analytics, AI, or real-time transaction processing need the right equipment in their data centers. Running workloads on servers with NVMe RAID is one approach to increasing performance, but organizations should also make sure they have the right servers and management software in place. For organizations looking to upgrade to meet business requirements, we researched the qualitative benefits of upgrading. We broke our evaluation down into three areas:

- PowerEdge R760 servers
- PERC 12
- Dell Technologies tools and services
- KIOXIA® CM6 Series Enterprise NVMe SSDs

PowerEdge R760 Servers

For fast and accurate performance, a PowerEdge R760 server makes sense as an upgrade over current systems, especially for workloads like AI and analytics. The PowerEdge R760 server is a 2U server that ships with up to two 4th Gen Intel® Xeon® Scalable processors and with up to 56 cores per socket. It keeps in-memory workloads humming at up to 4,800 megatransfers per second (MT/s) with DDR5 memory. Its Dell Smart Flow chassis helps to air-cool the components. Compared to the Dell PowerEdge R750 server, the PowerEdge R760 server delivers:

- **Up to 2.9x greater** AI inferencing on 4th Gen Intel Xeon Scalable processors with Intel® Deep Learning Boost (Intel® DL Boost) and Intel® Advanced Matrix Extensions (Intel® AMX)[2]
- **Up to a 20% increase** in the number of virtual desktop infrastructure (VDI) users supported on one server[3]
- **Up to 50% more** SAP® Sales and Distribution users supported on one server[4]

Dell PERC 12

Our testing results show that PERC 12 can run transactions up to 190% faster than PERC 11 on a PowerEdge R760 server. But there are other benefits that organizations can realize by upgrading to the latest-generation PERC 12 from PERC 11, or even from the prior-generation PERC 10.

PERC 12 provides a RAID solution that's powerful and easy to manage. It's based on the Broadcom SAS4116W processor, a dual-core ARM® A15 1.6 GHz RAID-on-Chip (ROC) designed to offer high I/O performance for data-intensive applications. PERC 12 supports 24 Gbps Serial-Attached SCSI (SAS) drives, increased cache memory speeds of up to 3,200 MHz, and a 16-lane host bus type.

The PERC 12 controller is a tri-mode ROC device. Tri-mode means that a PERC 12 controller can support either NVMe or SAS/Serial ATA (SATA) interfaces—though not both—from a single controller. This feature eliminates the need to use different controllers for SAS- and NVMe-supported servers. The controller has a slimline connector (or a NearStack® connector) for both PCIe and SAS/NVMe interfaces.

Table 1 shows a full list of features in the PERC 12 H965i Front.

Table 1 | Dell™ PERC 12 H965i Front features

Dell™ PERC 12 H965i Front		
Feature	**Description**	
Device interface	16 PHY SAS/SATA/PCIe® interfaces	
Form factor	PCIe® adapter card (HHHL) and custom Front and MX form factors	
Connectors	Two Amphenol® SlimSAS® x8 (SFF-8654 Series) connectors	
Device support	12 NVM Express® (NVMe®) devices (x2 connections) 24 SAS/SATA devices (limited by the number of controller back-end PHY slots)[5]	
Host bus type	16-lane, PCIe® Gen 4 compliant	
Data transfer rates	Up to 24 Gbps per PHY/port (SAS)[6]	
SAS controller	Dual-core ARM® A15 SAS4116W processor ROC	
Cache memory	2 MB shared L2 cache, 15 MB on-chip memory	
Key RAID and data-protection features	RAID levels 0, 1, 5, and 6, RAID spans 10, 50, and 60, online capacity expansion (OCE),[7] consistency check for background data integrity, physical disk power management (Dimmer Switch™), patrol read for media scanning and repair, 4K native sector support, self-encrypting drive (SED) support, and auto-resume after power loss during array rebuild	Load balancing, fast initialization for quick array setup, configurable stripe size, NVRAM "wipe" feature protects proprietary data once a card is decommissioned, up to 240 virtual drives,[5] DDF-compliant configuration on disk (COD), global and dedicated hot spare with revertible hot-spare, automatic rebuild, enclosure affinity, TRIM/UNMAP support for SSDs in pass-through mode
RAID management	PERC 12 card management applications include the Comprehensive Embedded Management (CEM), Dell™ OpenManage™ storage management, the Human Interface Infrastructure (HII) configuration utility, and the PERC command-line interface (CLI).	
Operating temperature	Maximum ambient temperature: 60°C	
Operating voltage	+3.3 V, +12 V, and +3.3 V Aux	
Optional SSD optimization	Hardware-accelerated I/O feature delivers high IOPS performance on SSD arrays; see User's Guide for details.	
Operating systems	Windows Server® with Hyper-V®, Red Hat® Enterprise Linux®, Ubuntu® Server, SUSE® Linux Enterprise Server (SLES), and VMware ESXi™ For specifications and interoperability details, see www.dell.com/ossupport.	

Broadcom® SAS4116W

The Broadcom SAS4116W high-performance I/O processor contributes to the higher performance of the Dell™ PERC 12.

The SAS4116W host interface supports 16-wide PCIe® lanes and complies with the PCIe 4.0 specification, offering up to 6.7 million random read IOPS and 1.1 million random write IOPS in RAID configurations.[8]

This processor is part of the eighth generation of SAS ROCs from Broadcom. It's based on the industry-leading Fusion-MPT™ architecture, and it features Tri-Mode SerDes technology.

Dell Technologies Tools and Services

Dell Technologies offers a wide range of software tools and services in its Dell™ OpenManage™ systems management portfolio. These tools and services help simplify management and support of PowerEdge servers, including the Dell PowerEdge R760 server. With the help of these services, staff can spend more time on value-add tasks and less time fixing problems. The result is reduced complexity and lower TCO. Some of these services include:

- **OpenManage Enterprise console.** The Dell OpenManage Enterprise console provides a comprehensive view of Dell Technologies servers and other devices on enterprise networks. It also supports plugins to automate the installation of firmware and drivers for streamlined updates.
- **OpenManage Server Administrator (OMSA).** OMSA provides tools for managing and monitoring RAID arrays. This includes the ability to create, modify, or delete RAID configurations, in addition to monitoring the status of RAID arrays and perform rebuilds.
- **Integrated Dell™ Remote Access Controller (iDRAC).** The OpenManage portfolio iDRAC, a remote server-management processor embedded within every PowerEdge server, can help IT administrators deploy, update, and monitor PowerEdge servers locally and remotely. This technology features a telemetry stream that IT admins can use to access sensor data from within the server. This data communicates parameters such as compute usage, aggregate temperature, and power consumption, which IT admins can use to proactively maintain servers and identify issues before they cause downtime.
  With iDRAC9, it's easy to manage the PERC 12 H965 card tested in our research. Without having to deploy an agent, IT admins can configure, deploy, update, and monitor the controller, either via the graphical user interface (GUI), with Redfish®, or through Dell's CLI, known as RACADM. With iDRAC9, customers can also perform real-time storage operations.
  iDRAC9 also supports Secured Component Verification (SCV), a supply-chain assurance offering. SCV enables customers to verify that the PowerEdge servers they receive from the factory match what was manufactured, and that nothing changed along the way.

Customers can get more out of their servers by taking full advantage of Dell Technologies' tools and services. They can also simplify management and support, which can help reduce TCO.

KIOXIA® CM6 Series Enterprise NVMe SSDs

Dell PowerEdge servers can significantly enhance their Sanitize Instant Erase (SIE) and compliance with Federal Information Processing Standards (FIPS) by using KIOXIA CM6 Series Enterprise NVMe SSDs. The SIE and FIPS security features help ensure data protection and meet stringent security requirements.

Achieving the Performance Edge

PowerEdge R760 servers with NVMe RAID—either PERC 11 or PERC 12—can help meet the performance requirements for workloads such as analytics, AI, and real-time processing. Our testing demonstrated that running transactions on the PowerEdge R760 server with PERC 12 instead of PERC 11 can result in an increase in storage performance of more than 10.8x. PERC 12 also reduced RAID rebuild times and SQL Server database restore times. This data suggests that an upgrade to PERC 12, built on Broadcom SAS4116W silicon, can help organizations achieve a performance edge and higher availability.

Our research also showed us that the PowerEdge R760 server should be considered in upgrade discussions because of its generational performance improvements over the PowerEdge R750 server. Additionally, Dell Technologies' tools and services offer a large portfolio of solutions that can reduce complexity and lower TCO.

Learn more about Dell PERC 12.

Learn more about the PowerEdge R760 servers.

Appendix

Our test configurations are shown in Table A1.

Table A1 | System configurations used for testing Dell™ PERC 12 versus PERC 11 on the Dell™ PowerEdge™ R760 server

Configuration	Single-Server Dell™ PERC 12	Single-Server Dell™ PERC 11
Server	1 x Dell™ PowerEdge™ R760	
Processor	Intel® Xeon® Platinum 8460Y+ processor	
Number of CPUs	2	
Cores/threads per CPU	32/64	
Cores/threads total	64/128	
Frequency (base/SCT/MCT)	2.2 GHz (max 4 GHz)	
Storage controller 01	Marvell Technology Group Ltd. Dell™ Boot-Optimized Server Storage (BOSS)-N1 Monolithic	
Disk	960 GB NVM Express® (NVMe®) Dell SK hynix® PE8010 RI M.2	
Number of disks	2	
Storage controller 02	Broadcom®/LSI Dell™ PERC H965i Front	Broadcom®/LSI Dell™ PERC H755N Front
Disk	1.6 TB NVMe® Dell Enterprise KIOXIA® Corporation CM6 MU (firmware: 2.2.0)	
Number of disks	8	
Storage controller 03	Broadcom®/LSI Dell™ PERC H965i Front	Broadcom®/LSI Dell™ PERC H755N Front
Disk	1.6 TB NVMe® Dell Enterprise KIOXIA® Corporation CM6 MU (firmware: 2.2.0)	
Number of disks	8	
Installed memory	256 GB	
Memory DIMM	SK hynix® HMCG78MEBRA174N Rank: Single-rank Memory type: error-correction code (ECC) DDR5	
Memory speed	4,800 megatransfers per second (MT/s)	
Number of memory DIMMs	16	
Network	10/25 gigabit Ethernet (GbE) remote direct memory access (RDMA) Broadcom® NetXtreme® E-Series BCM57508 quad-port Ethernet controller 2 x 100 GbE PCIe® Broadcom® NetXtreme® E-Series P2100D BCM57508 QSFP	
Operating system (OS)	Red Hat® Enterprise Linux® 9.2 (Plow)	
OS version	9.2	
OS kernel	Linux® 5.14.0-284.30.1.el9_2.x86_64	
Microsoft® SQL Server® version	Microsoft® SQL Server® 2022 (RTM-CU10) (KB5031778)—16.0.4095.4 (x64)	
BIOS version	1.5.6	1.5.6

[1] These test results obtained by Prowess Consulting using the fio benchmark confirm and extend similar testing conducted in 2023 by Tolly. For details on those test results, see: Tolly. "Dell Technologies Dell PowerEdge RAID Controller 12 (PERC 12) 16th Generation (16G) Server Performance vs PERC 11 & PERC 10." Commissioned by Dell Technologies. January 2023. **https://infohub.delltechnologies.com/en-US/section-assets/tolly223103delltechnologiespoweredgeraidcontroller12performance**.

[2] Testing commissioned by Dell Technologies in December 2022 and performed by Scalers AI demonstrated greater AI inferencing for object detection using INT8 on 4th Gen Intel® Xeon® Scalable processors for Dell™ PowerEdge™ R760 servers with 4th Gen Intel Xeon Scalable processors with updated Intel® Deep Learning Boost (Intel® DL Boost) instructions. Actual results will vary.

[3] Based on internal Dell Technologies testing results provided by Dell Technologies, in which a Dell™ PowerEdge™ R760 server with two Intel® Xeon® Gold 6454S processors hosted 220 VDI sessions, compared to a PowerEdge R750 server with two Intel Xeon Gold 6348 processors, which was able to host 183 VDI sessions.

[4] Based on test results approved by SAP under certification number 2023005, in which the Dell™ PowerEdge™ R760 server with two Intel® Xeon® Platinum 8480 processors with a total of 112 cores, 224 threads, and 2,048 GB DRAM hosted 72,250 SD users using SUSE® Linux® Enterprise Server (SLES) 15, compared to published results under certification number 2021026 for the PowerEdge R750 server with two Intel Xeon Platinum 8380 processors with a total of 80 cores, 160 threads, and 1,024 GB DRAM, which demonstrated support for up to 48,000 SD users using Red Hat® Enterprise Linux 8.2. Source: SAP. **SAP Standard Application Benchmarks** web page. Accessed August 2023.

[5] Silicon supports up to 240 virtual drives.

[6] SAS4 "22.5 Gbps" speed is used synonymously with "24G" and "24 Gbps" in some documents and applications; 22.5 Gbps is the data rate, whereas 24 Gbps is the link speed.

[7] OCE is supported from Dell™ PERC 12.1 onward.

[8] Broadcom. "**SAS4116W Tri-Mode RAID-on-Chip Product Brief**." October 2022.

The analysis in this document was done by Prowess Consulting and commissioned by Dell Technologies.

Results have been simulated and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

0424/220197