Upgrading to Dell PowerEdge R750 servers featuring Dell PowerEdge RAID Controllers (PERC 11): Stronger Apache Hadoop big data performance with RAID protection

The PowerEdge R750 solution boosted performance compared to a previous-generation PowerEdge R740 with PERC 10

Overview

Businesses of all sizes, and in all industries, rely on big data applications for everyday operations ranging from inventory management to risk mitigation to product recommendations. And the volume of important data continues to grow. With so much data to handle and critical business operations at risk, the server and storage solutions backing these applications must deliver both high performance and data protection.

Given the importance of big data, should organizations continue to rely on their older systems, or consider investing in new technologies? With careful tuning and configuration changes, organizations could possibly coax more performance from their legacy servers. But as multiple years pass and those servers grow older, IT teams inevitably run up against the limits of their aging hardware. New servers offer increased capabilities that can help deliver better performance and improved data protection for critical workloads.

The new Dell™ PowerEdge™ R750 server with the Broadcom®-based Dell PowerEdge RAID Controller 11 (PERC 11) offers both the high performance of non-volatile memory express (NVMe®) storage and the data redundancy of RAID. To quantify the performance advantages of this newer solution against a previous-generation PowerEdge R740 server with Dell PERC 10, we used a disk-intensive TeraSort workload to measure the performance of an Apache™ Hadoop® big data workload on both servers. Both solutions provided some data protection and redundancy via RAID, but only the newer Dell PowerEdge R750 server did so with NVMe drives. The newer server also delivered significant performance advantages, finishing the Apache Hadoop workload in less time and with greater throughput than the legacy system.

13.9%
less time to complete a TeraSort disk-intensive workload

16.3%
higher throughput on a TeraSort disk-intensive workload

How the explosion of data has driven the development of distributed databases

Traditional relational databases have been around for nearly half a century and have improved in countless ways. However, the underlying paradigm of how relational systems model data has remained largely consistent, with machines often needing to "scale-up," or become faster via hardware upgrades, to improve performance. In the last few decades, with the tidal waves of data brought on by internet, mobile, IoT, and other technologies, new clustered and distributed "scale-out" database systems have emerged with the goal of processing expansive amounts of data, both structured (as in relational systems) and unstructured (such as documents, pictures, text, etc.).

Hadoop is one such distributed system, comprising the MapReduce engine, the Hadoop Distributed File System (HDFS), Name Nodes, and Data Nodes. Clusters can be quite large in production, but the key to HDFS is its ability to break apart a very large data problem for processing.

According to the Apache wiki, organizations using Apache Hadoop include eBay, Facebook, Hulu, Spotify, Twitter, and dozens of smaller companies and educational institutions. Applications range from reporting/analytics and machine learning to search optimization to matching dating profiles to content generation and data aggregation.[1]

Quantifying the performance improvements of the Dell PowerEdge R750 with PERC 11

To assess the performance benefits of running Apache Hadoop on new Dell PowerEdge R750 servers, we deployed two small, virtualized Hadoop environments, one on the PowerEdge R750 and one on the PowerEdge R740. The PowerEdge R750 server had a Dell PERC 11 RAID controller, while the older PowerEdge R740 had a prior-generation Dell PERC 10 RAID controller. Both servers ran VMware® ESXi™ and used Linux® CentOS VMs and SSDs. The servers each contained six of the fastest drives that their controllers could support while providing RAID protection: NVMe SSDs for the Dell PERC 11 RAID controller, but only SAS SSDs for the Dell PERC 10 RAID controller.

Table 1: System configurations we used in performance testing. Source: Principled Technologies.

Server configuration information	Dell PowerEdge R750	Dell PowerEdge R740
Hardware		
Processors	2x Intel Xeon Gold 6348 28 cores each, 2.6 GHz	2x Intel Xeon Gold 6238R 28 cores each, 2.2 GHz
Storage controller	PERC H755N Front, 8GB cache	PERC H740P, 8GB cache
Disks	6x 1.6TB Dell Ent NVMe v2	6x WDC WUSTR6416BSS200
Total memory in system (GB)	256	256
Operation system name and version/build number	VMware ESXi, 7.0.3, 20036589	VMware ESXi, 7.0.3, 20036589
Software		
VM operating system	CentOS	
Benchmarking tools		
Hadoop big data performance	TeraSort benchmark, part of the HiBench suite of benchmarks	

We set up and configured the two servers remotely in a Dell lab. In configuring the servers for testing, we attached six SSD drives to the RAID controller on each server and created a RAID 5 logical drive on each system. We deployed a Hadoop cluster with one manager node and four workers on each server, and we configured the RAID 5 logical drive as storage for HDFS.

To test the two servers, we performed the following steps:

1. Rack and cable servers, verify BIOS and firmware levels on each server, and install vCenter Server.
2. Configure a 2-disk RAID 1 volume and a 6-disk RAID 5 volume on the servers.
3. Install VMware vSphere, and format datastores on each server.
4. Create VMs running CentOS on each server.
5. Install Apache Hadoop and Spark on the CentOS VMs, and create manager and worker nodes.
6. Load and configure HiBench suite with TeraSort dataset, and run tests.

For a detailed step-by-step methodology, see the science behind the report.

Comparing the two solutions' Apache Hadoop performance

To measure big data performance, we employed the TeraSort workload from the HiBench suite. In this workload, the TeraGen function generates input data, the TeraSort function uses MapReduce for sorting, and the TeraValidate function validates the output of the sorted data.[2] We selected this tool because it provides insight into the performance of Hadoop clusters and stresses the storage subsystem. The goal of testing was to generate performance data showing both run time and throughput on each platform. We ran the TeraSort workload three times and report the median of three runs. During testing, our experts relied on other performance data to confirm that the two platforms were functioning as we expected them to and that the configurations were comparable.

Figure 1 compares how long the two solutions took to complete the TeraSort workload. The Dell PowerEdge R750 server with Dell PERC 11 RAID controller required 4 minutes and 13 seconds—13.9 percent less time than the older Dell solution, which took 4 minutes and 54 seconds. This increased speed can translate to significant business benefits, especially at scale. Hundreds of thousands of data points are not useful on their own—they only become valuable after you've processed them and seen the patterns, trends, and revelations they contain. By completing these workloads faster with the PowerEdge R750, you gain more up-to-date insights to inform your next business decision.

Time to complete TeraSort workload

Min:sec | Lower is better

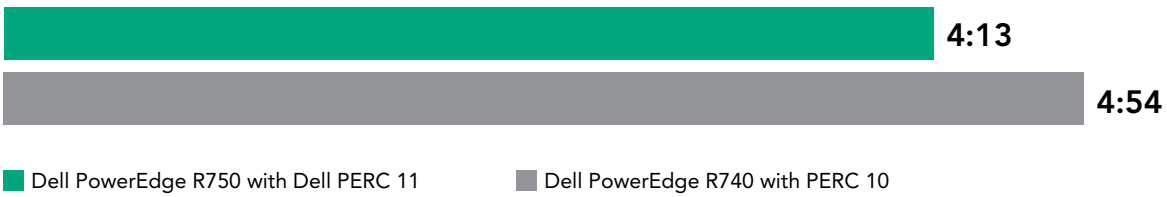

Figure 1: Time required for both solutions to complete a TeraSort workload on Apache Hadoop. Lower is better. Source: Principled Technologies.

Figure 2 shows the throughput rate each solution delivered while completing the TeraSort workload. This is the amount of data each solution processed in a given amount of time (in this case, the number of gigabytes in a second). The PowerEdge R750 server with PERC 11 RAID controller delivered 16.3 percent more gigabytes per second (GB/s) than the older solution. With higher throughput, the PowerEdge R750 could help organizations process more data in the same amount of time. Consider what this means for a business upgrading from the PowerEdge R740: In the span of a single day, that business could get more value from the same amount of data or gain the ability to use more data points in their everyday decision-making. As companies continue to collect more and more data, being able to put that data to work more quickly is a compelling proposition.

Throughput on a TeraSort workload
GB/s | Higher is better

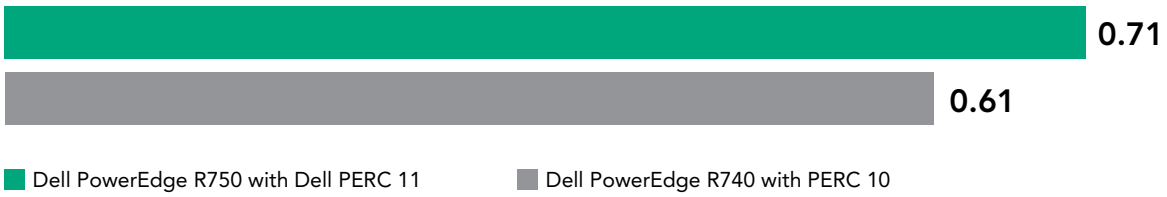

Figure 2: Throughput for both solutions during a TeraSort workload. Higher is better. Source: Principled Technologies.

About the Broadcom RAID-on-chip (ROC), SAS3916 chipset

The PERC processor in the Dell PERC 11 is a Broadcom ROC SAS3916 chipset. Broadcom based this chip on its Fusion-MPT architecture and says the chip "delivers enhanced performance and power reductions over previous generations. The ROC features Tri-Mode SerDes technology that enables a seamless operation of SAS, SATA or NVMe storage devices from any system design.

"The 16-port Tri-Mode ROC device provides SAS data transfer rates of 12, 6 and 3Gb/s per lane and SATA data transfer rates of 6 and 3 Gb/s per lane. The high-port count ROC helps eliminate storage bottlenecks with eight PCI Express® lanes and complies with the PCIe 4.0 specification, offering up to 3 million IOPS (JBOD mode) and up to 2.4 million IOPS in RAID (random reads)."[3]

About the Dell PowerEdge R750 server

The Dell PowerEdge R750 is a full-featured, general-purpose 2U rack server featuring 3rd Gen Intel® Xeon® Scalable processors. According to Dell, the PowerEdge R750 is purpose-built to optimize application performance and acceleration with PCIe Gen 4 compatibility, eight channels of memory per CPU, and up to 24 NVMe drives.[4] It also includes "I/O bandwidth and storage to address data requirements – ideal for: traditional corporate IT, database and analytics, virtual desktop infrastructure, AI/ML, and HPC."[5]

To learn more about the Dell PowerEdge R750, check out the spec sheet at https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-R750-spec-sheet.pdf.

How the Dell PowerEdge R750 with Dell PERC 11, using Broadcom technologies, lets you choose both performance and redundancy

Every year brings new technological innovations. In the two decades that we've been testing technology solutions here at Principled Technologies, we've seen dramatic improvements in processing, storage, and networking technologies, among others. Dell PowerEdge servers are no exception. In our work with PowerEdge servers from the ninth generation up to the latest, we've proven stronger performance generation over generation with a range of workloads.

In this testing, the Dell PowerEdge R750 with Broadcom technology-based Dell PowerEdge RAID Controller 11 offered clear performance advantages compared to its predecessor with PERC 10. As with any system upgrade, these advantages are likely due to multiple factors. We can attribute the strong advantage of the PowerEdge R750 solution in this test scenario to its use of newer processors, NVMe storage, and the updated storage controller with greater amounts of cache.

The Dell PERC 11, using the Broadcom RAID-on-chip (ROC) SAS3916 chipset, is especially powerful because it is compatible with a range of storage technologies: serial-attached SCSI (SAS), serial advanced technology attachment (SATA), and NVMe. The emergence of NVMe storage has brought significant storage performance gain, but until more recently, hardware-based RAID controllers have been compatible only with older, slower storage protocols such as SAS and SATA. Early in the NVMe transition, those who wanted to use RAID NVMe disks on servers had limited options to ensure that their data would be reliably available; the main options were software RAID or software-defined storage solutions. Thus, customers have had to sometimes prioritize either performance or redundancy.

Companies seeking the best of both worlds can now take advantage of new RAID controllers, such as the Dell PowerEdge RAID controllers (PERC 11), that are compatible with the NVMe protocol. According to Dell, Broadcom technology-based PERC 11 RAID controllers offer support for PCIe Gen 4, support for hot-swapping devices, non-volatile cache, secure enterprise key manager security, and more.[6] The Dell-Broadcom partnership on the RAID controller brings flexibility and performance together in the storage subsystem, allowing for system and application architects to use RAID 0, 1, 5, 6, 10, 50, and 60. This gives database architects the power to make new decisions on file placement and redundancy without sacrificing performance.

Conclusion

Companies increasingly turn to big data applications to solve business problems such as understanding customer habits and behavior, maintaining electronic health records, and detecting fraud. In a 2022 survey of executives, 97.0 percent had invested in big data initiatives, and 73.7 percent said their organizations had appointed a Chief Data Officer (CDO), up from 12 percent in 2012.[7]

With so much data at play, storage becomes an essential consideration for companies as they select hardware platforms to run their vital workloads. Two primary requirements for storage are fast performance and redundancy. Companies seek storage that can quickly put actionable insights into the hands of decision makers. At the same time, it is a fact of life that storage media occasionally fails, and no company wants to risk the potentially very large expense of losing vital business data.

In our testing with a disk-intensive TeraSort big data workload, we found that a current-generation Dell PowerEdge R750 server with a Broadcom-based Dell PERC 11 RAID controller with NVMe SSDs outperformed a previous-generation PowerEdge R740 server with a Dell PERC 10 RAID controller. The newer PowerEdge R750 solution completed the workload in 13.9 percent less time and with 16.3 percent better throughput, indicating that it could speed the delivery of insights to decision makers. With its PERC 11 RAID controller, it also offers RAID support for today's speedy NVMe SSDs. Data-driven organizations must weigh the cost savings of putting off purchases of newer servers against the benefits of investing in those same newer, more powerful servers with RAID support for NVMe drives. Our testing shows that with the Dell PowerEdge R750 server with a Broadcom technology-based Dell PERC 11 RAID controller, buyers need not choose between redundancy and high performance.

1. Apache, "Powered by Apache Hadoop," accessed July 19, 2023, https://cwiki.apache.org/confluence/display/HADOOP2/PoweredBy.
2. Ahmed, N., Barczak, A.L.C., Susnjak, T. et al. "A comprehensive performance analysis of Apache Hadoop and Apache Spark for large scale data sets using HiBench," Journal of Big Data, accessed July 13, 2023, https://doi.org/10.1186/s40537-020-00388-5.
3. Broadcom, "SAS3916 12Gb/s SAS Tri-Mode RAID-on-Chip (ROC)," accessed July 13, 2023, https://www.broadcom.com/products/storage/raid-on-chip/sas-3916.
4. Dell, "Dell EMC PowerEdge R750 Spec Sheet," accessed July 13, 2022, https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-R750-spec-sheet.pdf.
5. Dell, "Dell EMC PowerEdge R750 Spec Sheet."
6. Dell, "Dell PowerEdge RAID Controller 11 User's Guide PERC H755 adapter, H755 front SAS, H755N front NVMe, H755 MX adapter, H750 adapter SAS, H355 adapter SAS, H355 front SAS, H350 adapter SAS, H350 Mini Mo," accessed July 13, 2023, https://www.dell.com/support/manuals/en-pa/perc-h755/perc11_ug/technical-specifications-of-perc-11-cards?guid=guid-aaaf8b59-903f-49c1-8832-f3997d125edf&lang=en-pa.
7. New Vantage Partners, "Data and AI Leadership Executive Survey 2022, Executive Summary of Findings," accessed July 13, 2022, https://c6abb8db-514c-4f5b-b5a1-fc710f1e464e.filesusr.com/ugd/e5361a_2f859f3457f-24cff9b2f8a2bf54f82b7.pdf.

Read the science behind this report at https://facts.pt/mJepPM7 ▶

Principled Technologies®

Facts matter.®

This project was commissioned by Dell Technologies.