Technical Research Report

PROWESS

A Comparative Study of Hardware and Software RAID Performance and Benefits

Prowess Consulting compared hardware and software RAID arrays on Dell™ PowerEdge™ servers across a variety of scenarios to determine which performed best in each usage.

Executive Summary

Hardware RAID remains the preferred choice for most enterprise workloads due to its ability to deliver lower latency, higher input/output operations per second (IOPS), and superior rebuild performance under load, compared to software RAID solutions. Using a dedicated RAID controller, hardware RAID helps ensure stable, predictable performance, particularly as disk configurations scale, making it indispensable for mission-critical environments where consistency and rapid recovery are paramount. While hardware RAID excels in these areas, software RAID has gained traction in smaller setups. By making use of system resources, software RAID can deliver impressive bandwidth and performance in configurations with fewer disks, particularly when paired with NVM Express® (NVMe®) drives. Moreover, software RAID can provide much higher bandwidth than hardware RAID.

Highlights of the Test Results

In Prowess Consulting's testing, hardware-based and software-based RAID excelled in different areas:

>>

Up to **281%** higher median IOPS for **hardware RAID**

Up to **99%** lower median latency for **hardware RAID**

Up to **3,521%** higher median bandwidth for **software RAID**

Prowess Consulting investigated the benchmarked performance of software RAID versus hardware RAID to determine in which circumstances each performed better. This report, sponsored by Dell Technologies, presents detailed results from various test scenarios that we conducted on four similarly configured Dell™ PowerEdge™ R760 servers using software and hardware RAID arrays:

1. One PowerEdge R760 server with the Red Hat® Enterprise Linux® operating system (OS) with built-in multiple-device RAID (software RAID)
2. One PowerEdge R760 server with the Windows Server® 2022 OS with Dell™ PowerEdge RAID Controller (PERC) S160 software RAID (software RAID)
3. One PowerEdge R760 server with the Red Hat Enterprise Linux OS with Dell PERC H965i Front (hardware RAID)
4. One PowerEdge R760 server with the Windows Server 2022 OS with Dell PERC H965i Front (hardware RAID)

For each system, we tested a series of drive size configurations:

- RAID 0 with 1, 2, 4, and 8 drives
- RAID 10 with 4 and 8 drives
- RAID 5 with 4 and 6 drives

We measured each system's performance using benchmarking tools, and we compared the systems in terms of key metrics such as IOPS, bandwidth, latency, and rebuild times to determine the following:

- Whether software RAID delivers higher bandwidth due to full access to an NVMe drive (x4 PCIe®)
- Under what circumstances hardware RAID performs better than software RAID, and vice versa

Modern Workloads Place More Demands on Storage

As modern workloads push the limits of storage performance, hardware RAID remains the preferred choice for most demanding applications. Using a dedicated RAID controller with hardware RAID provides lower latency and higher IOPS, especially in multi-disk setups. These characteristics deliver predictable, high-performance input/output (I/O) suited to transactional databases, virtualized environments, and enterprise applications that depend on consistency and speed. Using an independent controller with hardware RAID also offloads I/O tasks from the host CPU. Doing so allows for better scalability while minimizing system resource utilization.

However, software RAID has gained traction in recent years, particularly for more flexible deployments. Often included as part of operating systems like Linux and Windows Server, software RAID offers unique advantages, such as hardware independence and straightforward disk migration across systems. Additionally, software RAID can make use of the high bandwidth of NVMe drives via x4 PCIe channels. This capability makes it advantageous for bandwidth-intensive applications that are less latency-sensitive.

To understand the current benefits of software and hardware RAID and how they compare, Prowess Consulting conducted a research study that tested four similarly configured PowerEdge R760 servers using the two different RAID types. Through this study, we aimed to provide a deeper understanding of the performance and resource utilization of both RAID types and the scenarios in which one type is better than the other.

Benchmarking Tools

The benchmarking tools used in our testing included the following:

- **Flexible I/O Tester (fio),** a free and open-source tool that is used for benchmarking and testing the performance of storage systems. Fio can be used to generate a wide range of I/O workloads to simulate various real-world scenarios.
- **Iometer,** an I/O subsystem measurement and characterization tool for single and clustered systems. Iometer was originally developed by Intel and is now maintained by an international group of individuals.
- **Performance Monitor (PerfMon),** a system monitoring tool for Windows® that monitors computer activities such as CPU usage and memory usage and that also measures IOPS. PerfMon performs asynchronous I/O and allows the configuration of disk parameters such as maximum disk size, starting disk sector, and number of outstanding I/O operations.
- **Input/output statistics (iostat),** which collects and shows OS storage I/O statistics. Iostat is often used to identify performance issues with storage devices, including local disks or remote disks accessed over file systems such as Network File System (NFS).
- **HammerDB,** an open-source benchmarking tool designed to evaluate database performance by simulating transactional and analytical workloads. HammerDB is widely used to measure and compare the performance of database systems under various conditions, providing metrics such as new orders per minute (NOPM) to assess transactional throughput.

When to Use Which Type of RAID Controller

Our testing underscores that hardware RAID and software RAID have different uses for which they are best suited. An important part of our testing was not to illustrate the general superiority of either hardware RAID or software RAID, but to discern the types of use cases that can be best served by each.

When hardware RAID is more desirable:	When software RAID is more desirable:
1. **High-performance, latency-sensitive applications:** • **Database servers:** If you're running a transactional database (such as Microsoft® SQL Server® or Oracle® Database) that demands low latency for real-time data processing, hardware RAID is ideal. Its dedicated RAID controller reduces I/O bottlenecks, resulting in more predictable, lower latency, which is critical for high-frequency read/write operations. 2. **Data integrity and redundancy (RAID 5 or 10):** • **Enterprise environments:** Hardware RAID is typically more reliable in complex RAID configurations (like RAID 5), where parity calculations can be offloaded to the RAID controller, minimizing the impact on the system's main CPU. This is especially important in mission-critical environments where data redundancy and performance must be balanced, such as in large-scale enterprise storage systems. • **High-availability setups:** For systems that require failover and redundancy (such as those in financial services and healthcare), hardware RAID provides advanced features like battery-backed cache and instant failover, helping ensure data integrity during outages or power loss. 3. **Better scaling with multiple drives:** • **Storage arrays with many drives:** As shown in our test data, hardware RAID scales more effectively when multiple disks are used. This makes it preferable in scenarios with high disk counts, such as large storage pools, media servers, or environments where IOPS scalability is crucial. • **RAID 10 or higher configurations:** Hardware RAID typically handles complex striping and mirroring setups (such as RAID 10 and RAID 50) with greater efficiency, providing high fault tolerance and performance.	1. **Flexible storage management and customization:** • **Containerized environments:** In modern DevOps environments where storage might be managed via container orchestration systems (such as Kubernetes®), software RAID can offer better integration, flexibility, and automation than hardware RAID. 2. **CPU- and memory-rich systems:** • **NVMe® setups with high bandwidth:** Software RAID can exploit high-bandwidth NVMe disks by making use of the host CPU. If latency isn't critical and the system has a powerful CPU, software RAID can provide higher throughput for data-intensive applications where massive amounts of sequential data are read or written. • **Non-latency-critical workloads:** For workloads where latency isn't a primary concern (such as cold storage or backup servers), software RAID can be a viable option without significant performance penalties. 3. **Cloud environments:** • **Cloud-native environments:** Software RAID can work better in cloud environments where dedicated RAID hardware is impractical or unavailable. • **Backup and archival solutions:** If the system's primary function is long-term storage of data with minimal access (like backup or archival), the performance benefits of hardware RAID might be unnecessary. Software RAID can provide sufficient redundancy without the added cost of a RAID controller. 4. **Workloads tolerant to latency:** • **Media storage or archival:** Software RAID could be used in media archiving or content distribution where high bandwidth is required but latency isn't a critical issue.

Hardware-based RAID also has some unique benefits:

- **Vigorous qualification and validation testing:** Hardware RAID solutions undergo rigorous qualification and validation testing by manufacturers to ensure high reliability and performance under various workloads and failure scenarios. These tests include compatibility checks with diverse operating systems, storage drives, and server environments. They also include stress tests to simulate real-world conditions. This extensive vetting process reduces the likelihood of unexpected failures and ensures consistent performance. As a result, hardware RAID is a dependable choice for mission-critical applications where data integrity and availability are paramount.
- **Advanced management software, both CLI and GUI:** Hardware RAID controllers are often complemented by sophisticated management software, offering both command-line interface (CLI) and graphical user interface (GUI) options for configuration and monitoring. These tools simplify RAID setup, management, and troubleshooting. They also provide intuitive dashboards, real-time health metrics, and automated alerts for disk failures or performance issues. In contrast, software RAID solutions like MDRAID can require intricate CLI commands, which can make management more cumbersome for administrators. The user-friendly software accompanying hardware RAID helps reduce administrative overhead and ensures faster issue resolution.

- **Cache for low-latency accesses:** Hardware RAID controllers typically include dedicated cache memory to accelerate read and write operations. By buffering frequently accessed or write-pending data, these caches significantly reduce latency and improve overall performance. This is especially true for workloads with high IOPS demands. Additionally, features like battery-backed or non-volatile memory ensure data in the cache is preserved during power outages, enhancing both speed and data protection. This level of caching optimization is typically unavailable or less effective in software RAID implementations.

Test Results and Analysis

In considering the test results from the perspective of a hypothetical company considering the capabilities of hardware and software RAID, we first reviewed the performance of each system using the selected benchmark tests. In general, hardware RAID provided higher performance and lower latency while software RAID supplied higher bandwidth.

Hardware RAID Wins for IOPS

IOPS generally increases as disk quantity rises because additional drives allow more parallel processing of I/O operations. However, this increase is not strictly linear. The rate of improvement can diminish with higher disk counts. Hardware RAID consistently demonstrates higher IOPS than software RAID across multi-disk setups. For configurations with more than one disk, hardware RAID's dedicated controller provides a reliable performance advantage. In our testing, hardware RAID yielded 8–281% higher median IOPS over software RAID at 2, 4, 6, and 8 disks, which underscores its suitability for workloads that require both high IOPS and predictability (see Table 1). This reliability is especially valuable in environments where performance must scale predictably with additional disks.

In contrast, software RAID can leverage the CPU's processing power to handle higher IOPS. This is especially true when fewer disks are used or when system resources (CPU and memory) are sufficient to manage the additional I/O load. With a single disk, software RAID shows 30% higher IOPS than hardware RAID, demonstrating that it can excel in scenarios with lighter workloads or where cost efficiency and bandwidth are prioritized over latency (Table 1). However, as disk quantity increases, software RAID's dependence on the system's CPU can become a bottleneck, which can result in lower IOPS than hardware RAID beyond the single-disk configuration. These results suggest that while software RAID can be advantageous in small-scale deployments, hardware RAID generally offers more consistent IOPS at higher disk counts.

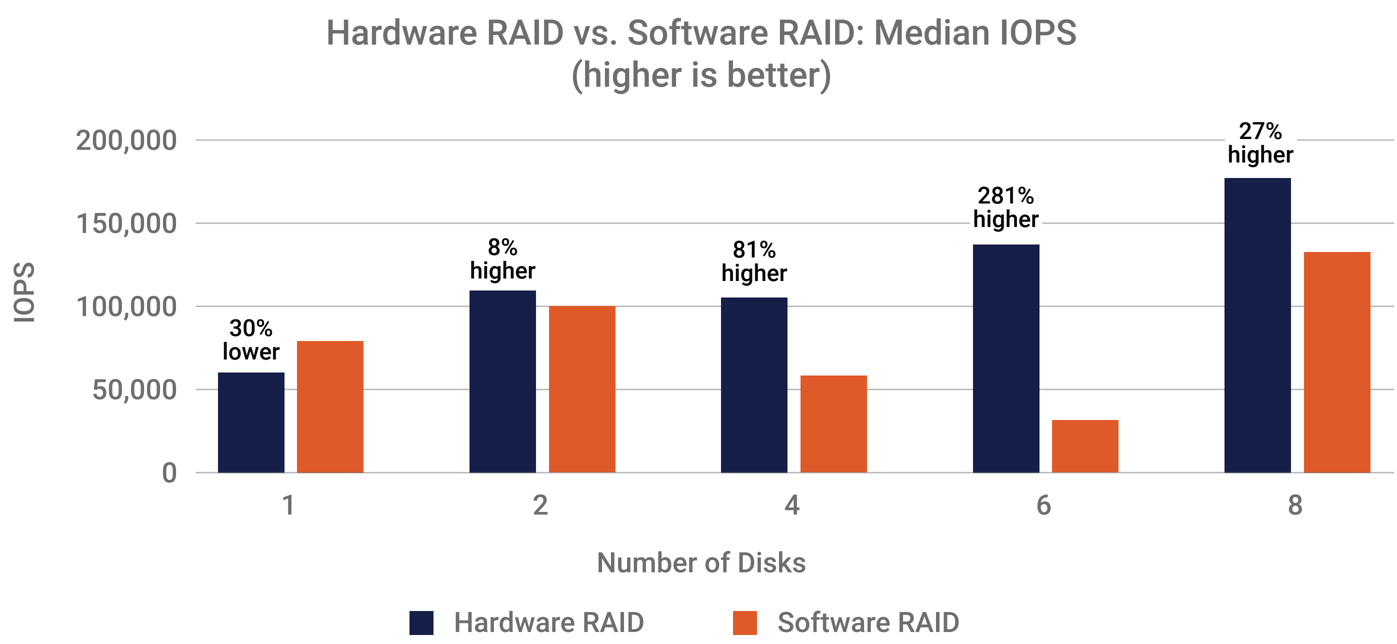

Figure 1 | Median IOPS results by number of disks for hardware RAID and software RAID across RAID 0, 5, and 10

Table 1 | Median IOPS results by number of disks for hardware RAID and software RAID across RAID 0, 5, and 10

Disk Quantity	1	2	4	6	8
Hardware RAID Median IOPS	56,564	107,142	103,723	139,710	177,936
Software RAID Median IOPS	80,917	98,970	57,240	36,679	140,144
Hardware RAID Median IOPS Comparison	30% lower	8% higher	81% higher	281% higher	27% higher

Factors such as worker count (number of threads) and queue depth (number of concurrent I/O requests) also impact IOPS performance. Software RAID benefits from higher worker counts and queue depths by utilizing the system's full resources. Doing so can yield performance gains but also adds variability as CPU utilization rises. Hardware RAID, on the other hand, demonstrates more stable IOPS with less dependency on system resources, given its independent RAID controller.

Hardware RAID Outperforms Software RAID in Latency

Low latency is crucial for high-performance workloads. Hardware RAID consistently outperformed software RAID in this regard, providing significantly lower latency across all disk configurations. The independent RAID controller of hardware RAID ensures more predictable and stable latency by handling I/O operations efficiently without overburdening system resources. In our testing, hardware RAID maintained latencies under 600 microseconds for most configurations, with only the four-disk setup showing higher latency due to an anomaly in processing demands (see Table 2).

On the other hand, software RAID exhibits substantially higher latencies across disk quantities due to its dependence on the system's CPU and memory for processing. Even with a single disk, software RAID's latency was several orders of magnitude higher than hardware RAID's. Moreover, the gap persisted with additional disks, which reflects the variability and system overhead inherent to software RAID. Software RAID's median latency was 99% higher or more across multiple configurations, as it lacks the dedicated resources to manage I/O with the same efficiency as hardware RAID (see Table 2).

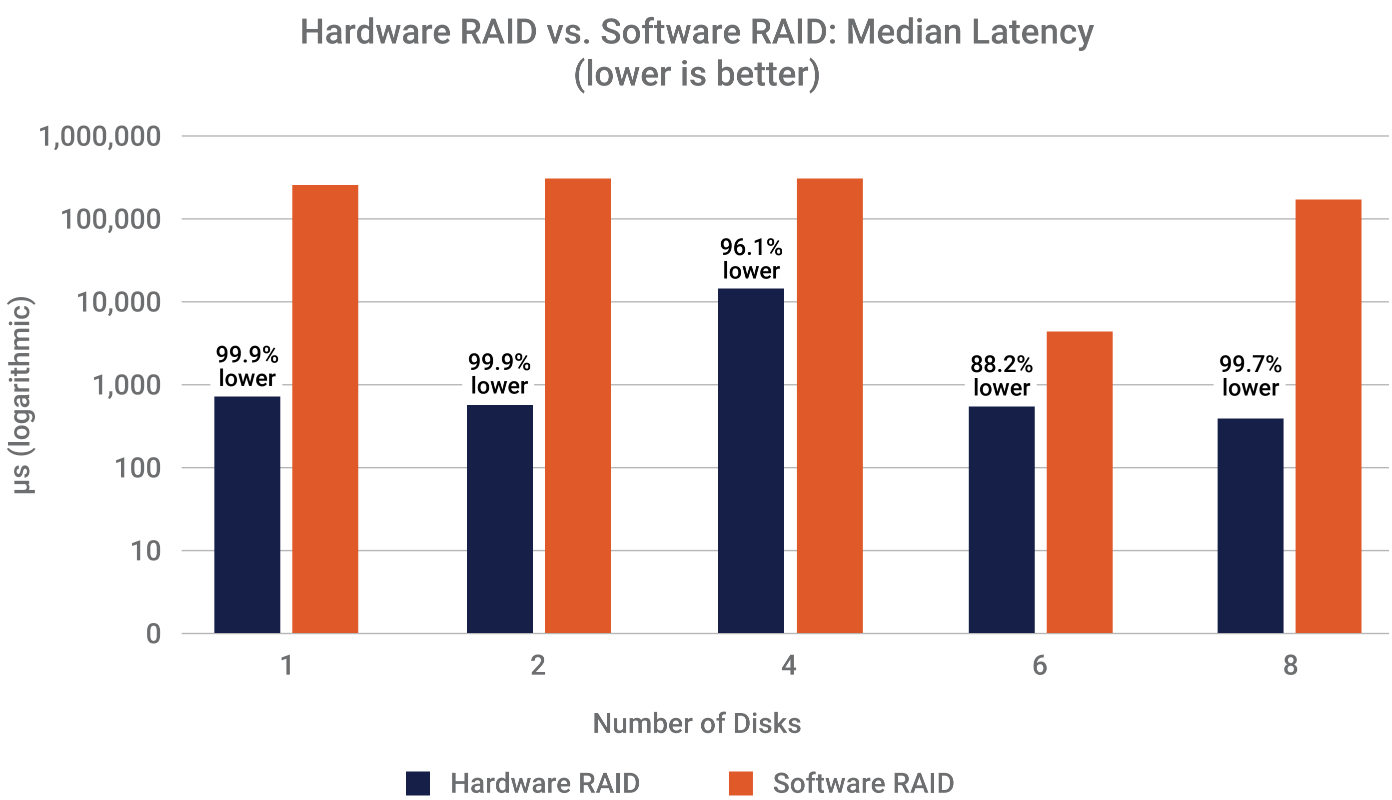

Figure 2 | Median latency results by number of disks for hardware RAID and software RAID across RAID 0, 5, and 10

Table 2 | Median latency results by number of disks for hardware RAID and software RAID across RAID 0, 5, and 10

Disk Quantity	1	2	4	6	8
Hardware RAID Median Latency (99.9%)	696 µs	591 µs	12,992 µs	597 µs	597 µs
Software RAID Median Latency (99.9%)	329,728 µs	372,736 µs	329,728 µs	5,056 µs	234,496 µs
Hardware RAID Median Latency Comparison	99.8% lower	99.8% lower	96.1% lower	88.2% lower	99.7% lower

Software RAID Is the Clear Winner for Bandwidth

While hardware RAID provides stable performance across various disk configurations, its bandwidth scalability can be limited by the fixed processing capacity of the RAID controller. This can be particularly evident as the number of disks increases. In our testing, hardware RAID exhibited predictable but lower bandwidth than software RAID, achieving its highest median bandwidth of 473,251 KB/s in the four-disk configuration (see Table 3). This consistency makes hardware RAID suitable for applications where stable, reliable throughput is prioritized over peak bandwidth, such as database systems or virtualized environments that benefit from steady I/O without unexpected spikes.

Conversely, software RAID capitalizes on the host system's CPU and memory resources. This functionality enables it to achieve significantly higher bandwidth in most configurations. Software RAID outperformed hardware RAID by a substantial margin in this metric. It achieved up to 3,521% higher bandwidth in the two-disk configuration and showed similar dominance in the one-, six-, and eight-disk setups (Table 3). This bandwidth advantage is particularly beneficial in data-intensive environments, such as large-scale content delivery networks (CDNs) or AI/machine learning (ML) workloads, where high throughput is more valuable than low latency. However, software RAID's bandwidth performance in the four-disk configuration was 37% lower than hardware RAID's, possibly due to bottlenecks in system resource allocation at that specific disk count.

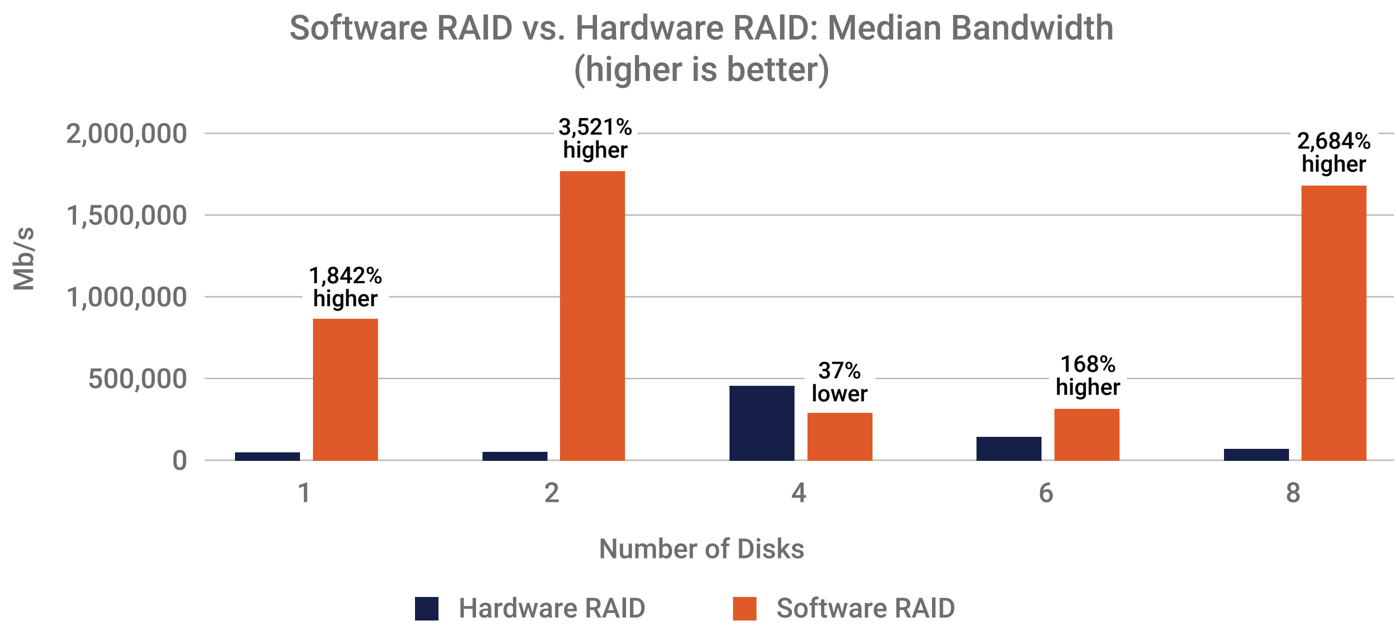

Figure 3 | Median bandwidth results by number of disks for hardware RAID and software RAID across RAID 0, 5, and 10

Table 3 | Median bandwidth results by number of disks for hardware RAID and software RAID across RAID 0, 5, and 10

Disk Quantity	1	2	4	6	8
Hardware RAID Median Bandwidth	46,706 Mb/s	49,666 Mb/s	473,251 Mb/s	109,560 Mb/s	59,303 Mb/s
Software RAID Median Bandwidth	906,808 Mb/s	1,798,222 Mb/s	299,063 Mb/s	293,115 Mb/s	1,650,948 Mb/s
Software RAID Median Bandwidth Comparison	1,842% higher	3,521% higher	37% lower	168% higher	2,684% higher

While software RAID can scale bandwidth more easily by making use of CPU power and PCIe lanes, its latency remains significantly higher, especially at larger disk counts. This makes software RAID better suited for applications that prioritize bandwidth over latency, such as sequential data reads/writes or scenarios where cost and flexibility are more important than rapid response times.

Software RAID's reliance on system resources also means that worker count and queue depth significantly impact its bandwidth capabilities. With more threads available for I/O requests, software RAID can achieve high throughput, though extreme concurrency levels might strain the CPU and memory. This reliance can cap its performance gains. Hardware RAID, while benefiting modestly from increased worker counts and queue depths, tends to exhibit diminishing returns as concurrency increases, which reflects the constraints of its RAID controller.

Performance Under Rebuild

Rebuild performance measures how quickly a RAID array can return to optimal performance after a drive failure or replacement. During the rebuild process, the RAID system regenerates lost data by using parity or mirroring information across the remaining drives. This functionality is critical for maintaining data integrity and availability, especially in enterprise environments. However, while a RAID array is rebuilding, overall system performance typically degrades; the array must balance the recovery process with servicing ongoing I/O requests.

High rebuild performance under load is thus essential to minimize the impact on production workloads and reduce the time a system remains in a vulnerable state. Fast rebuilds ensure that the RAID array can recover data quickly, which limits exposure to potential data loss should another drive failure occur during the rebuild process.

In our testing, we measured the rebuild performance under load by dividing the median IOPS during the rebuild by the number of minutes required for the RAID setup to return to optimal operation. The results revealed significant differences between hardware and software RAID, with hardware RAID exhibiting far superior rebuild performance (see Table 4).

Table 4 | Median rebuild performance under load results by number of disks for hardware RAID and software RAID across RAID 0, 5, and 10

Disk Quantity	4	6	8
Hardware	30,130.85 IOPS/min	2,318.12 IOPS/min	11,903,838.53 IOPS/min
Software	251.22 IOPS/min	91.55 IOPS/min	453.59 IOPS/min
Relative Performance	119x	25x	26,243x

These results indicate that hardware RAID consistently outperforms software RAID during rebuilds. With four disks, hardware RAID was 119 times more efficient in terms of rebuild performance compared to software RAID. This margin widened dramatically at a high disk count, with hardware RAID outperforming software RAID by a factor of 26,243 with eight disks.

Software RAID, on the other hand, displayed much lower rebuild performance under load, likely due to its reliance on system resources such as the CPU for handling both I/O and the rebuild process. This introduces a significant bottleneck during rebuild operations, especially as the number of disks increases.

Memory Usage: Optimal Conditions and Under Rebuild

Memory usage is a critical factor in RAID array performance. Different RAID levels demand varying amounts of system memory to manage data integrity and redundancy, especially during degraded or rebuild states. Hardware RAID offloads some of these memory-intensive tasks to a dedicated RAID controller, which has its own cache to accelerate data access and minimize impact on system memory. By contrast, software RAID relies on the system's CPU and memory for parity calculations, data reconstruction, and I/O management, which can increase memory consumption and reduce resources available for other applications. Effective memory utilization can thus significantly affect RAID performance, particularly in high-demand environments or with RAID levels like RAID 5 that involve complex parity calculations. Insufficient memory can lead to performance bottlenecks and increased latency.

Our testing showed that hardware RAID used 43–46% less memory than software RAID in both optimal and rebuild states for RAID 5 (see Figure 4 and Table 5).

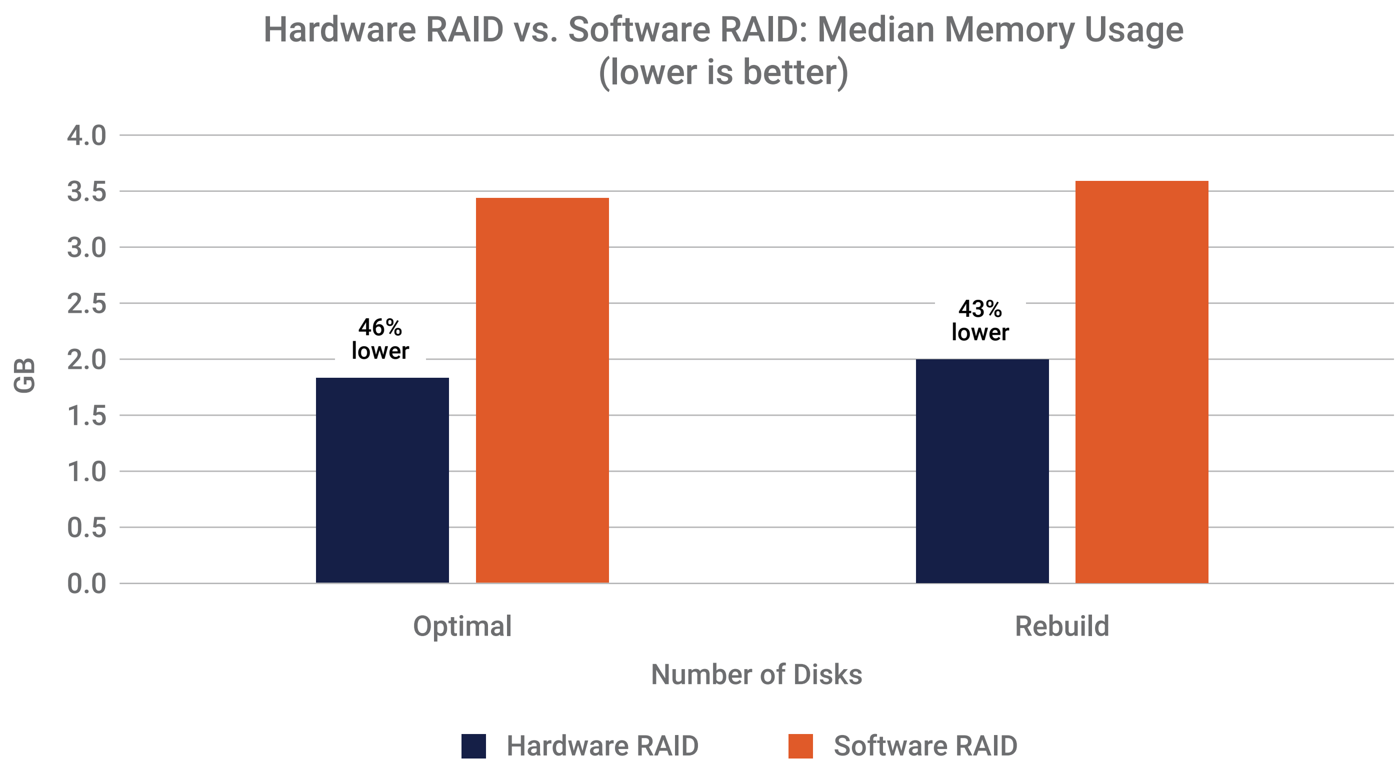

Figure 4 | Median memory usage for hardware RAID and software RAID in both optimal and rebuild states for RAID 5

Table 5 | Median memory usage for hardware RAID and software RAID in both optimal and rebuild states for RAID 5

RAID Operating State	Optimal	Rebuild
Hardware RAID	1.86 GB	2.01 GB
Software RAID	3.48 GB	3.56 GB
Relative Memory Usage for Hardware RAID	46% lower	43% lower

Real-World Application Testing

In order to assess the relative performance of hardware-based and software-based RAID configurations in real-world conditions, Prowess Consulting used HammerDB. HammerDB is a popular open-source benchmarking tool used to evaluate database performance by simulating transactional and analytical workloads. It uses warehouses as its units of data to simulate the scale of transactional workloads. Prowess Consulting used NOPM as the performance metric for this benchmark testing to focus specifically on the transactional throughput of the hardware and software RAID configurations.

We examined RAID 5 for this testing because it is common in production environments due to its blend of data protection and storage efficiency, and because it places additional strain on the system due to parity calculations. RAID 5 uses striping with parity and distributes data and parity information across all disks. This configuration provides fault tolerance by using parity data to reconstruct lost information if a single disk fails, but it allows more storage efficiency because only one disk's worth of space is used for parity.

Testing with HammerDB showed hardware-based RAID to be consistently several times faster than software-based RAID, ranging from 3.17x NOPM to 14.6x NOPM across optimal, degraded, and rebuild disk conditions (see Figures 5 and 6 and Tables 6 and 7). Prowess Consulting engineers simulated a failed disk (degraded) scenario by removing one of the disks; the rebuild scenario had the disk restored but measured performance as data and parity information was being calculated and copied to the "replacement" disk. This disparity in performance across disk conditions was most pronounced in the four-disk configuration as the number of warehouses increased (see Figure 5 and Table 6).

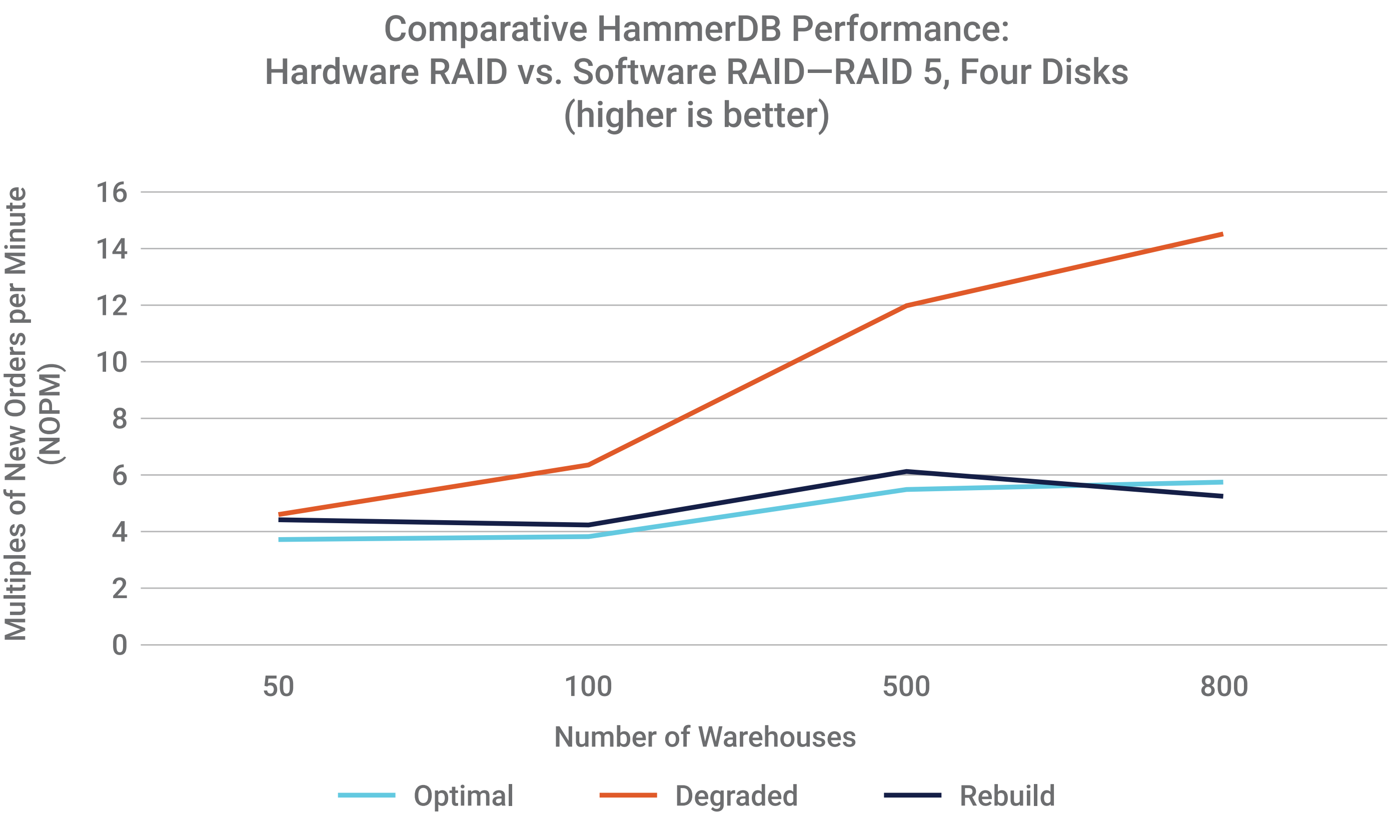

Figure 5 | Comparative HammerDB performance for hardware RAID and software RAID on RAID 5 with four disks for optimal, degraded, and rebuild disk conditions for 50, 100, 500, and 800 warehouses

Table 6 | Comparative HammerDB performance for hardware RAID and software RAID on RAID 5 with four disks for optimal, degraded, and rebuild disk conditions for 50, 100, 500, and 800 warehouses

Warehouses	50	100	500	800
Optimal	3.54x	3.60x	5.47x	5.70x
Degraded	4.69x	6.67x	11.91x	14.68x
Rebuild	4.57x	4.28x	6.06x	4.97x

The performance difference was more constant in the six-disk configuration, but hardware RAID performed even better against software RAID with more disks in the degraded and rebuild scenarios (see Figure 6 and Table 7).

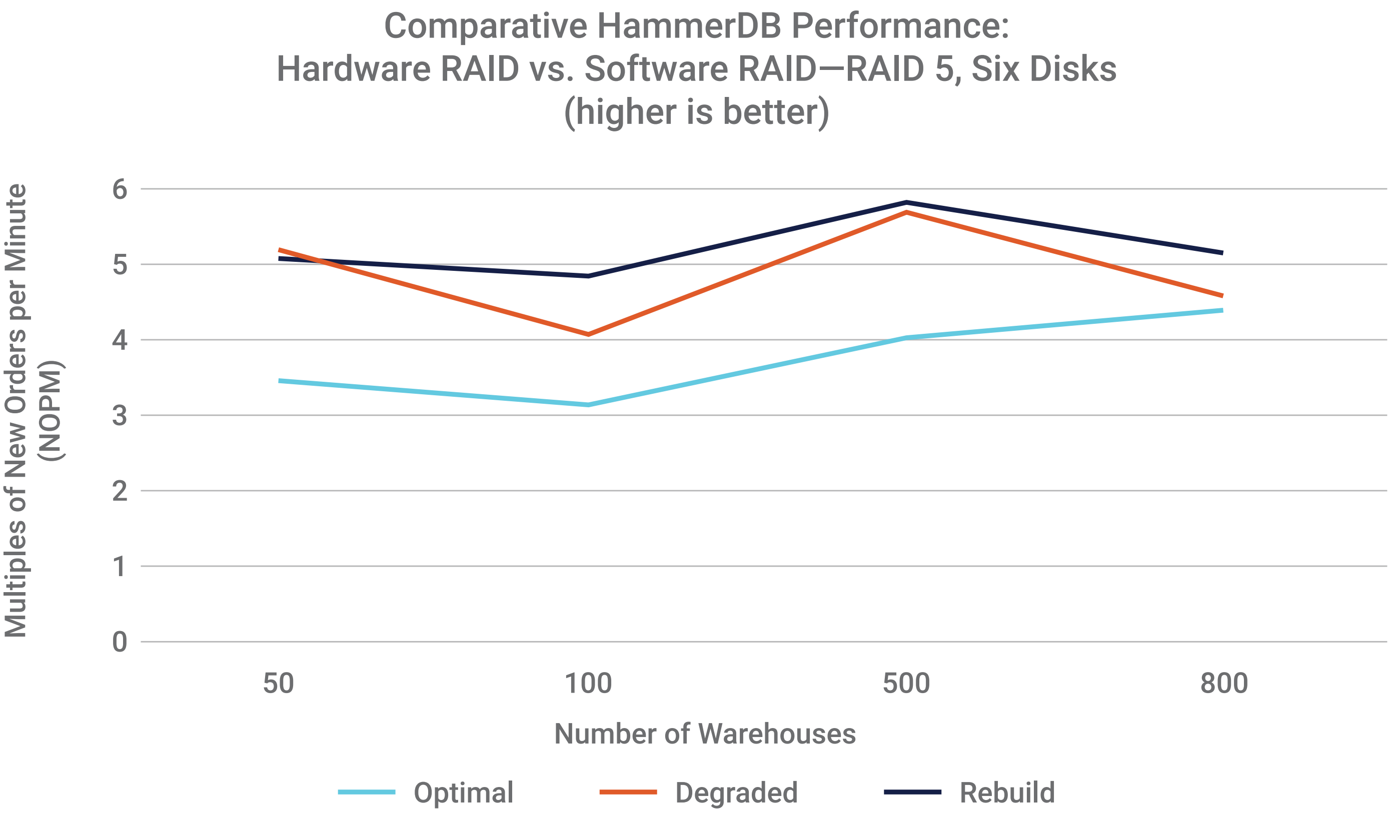

Figure 6 | Comparative HammerDB performance for hardware RAID and software RAID on RAID 5 with six disks for optimal, degraded, and rebuild disk conditions for 50, 100, 500, and 800 warehouses

Table 7 | Comparative HammerDB performance for hardware RAID and software RAID on RAID 5 with six disks for optimal, degraded, and rebuild disk conditions for 50, 100, 500, and 800 warehouses

Warehouses	50	100	500	800
Optimal	3.48x	3.17x	3.98x	4.41x
Degraded	5.22x	4.13x	5.56x	4.68x
Rebuild	5.13x	4.75x	5.71x	5.17x

In degraded conditions, hardware RAID typically outperforms software RAID by offloading parity calculations and data reconstruction to a dedicated RAID controller. This benefit can vary with the disk count due to how RAID 5 organizes and manages data and parity across disks.

RAID 5 uses parity data to rebuild lost data in degraded conditions. In the four-disk configuration, the amount of parity calculations is generally lower, meaning that the RAID controller in hardware RAID can complete parity calculations and serve data requests more quickly. In contrast, a six-disk RAID setup has a larger dataset and more parity information to handle. Our testing illustrated that this can dilute the performance advantage of hardware RAID relative to software RAID.

In addition, in a six-disk array, when one disk fails, data has to be reconstructed from more complex parity and data patterns spread across additional disks. This increases the read and write operations required to retrieve the full data, which is computationally more intensive, especially for software RAID. Hardware RAID mitigates this impact with its controller, but it might still lose some of its relative advantage due to the increased overhead. With six disks, the controller has to handle more data segments and parity blocks per data request. While hardware RAID remains more efficient, its relative advantage can decrease because it's working with a larger dataset.

Other Performance-Improving Features of PowerEdge R760 Servers

PowerEdge R760 servers have additional features that improve performance, especially for workloads like AI and analytics. The PowerEdge R760 server is a 2U server powered by Intel® Xeon® processors, which include built-in Intel® Accelerator Engines that boost workload performance, such as Intel® Deep Learning Boost (Intel® DL Boost) and Intel® Advanced Matrix Extensions (Intel® AMX).

Dell Technologies' tools and services help users get even more out of PowerEdge servers for better total cost of ownership (TCO):

- **Dell™ OpenManage™ systems management portfolio** offers a wide range of software tools and services that simplify the management and support of PowerEdge servers. These tools and services allow IT staff to spend more time on value-added tasks and less time fixing problems. The Dell OpenManage Enterprise console provides a comprehensive view of Dell™ servers and other devices on enterprise networks. OpenManage Server Administrator provides tools for managing and monitoring RAID arrays, including the ability to create, modify, or delete RAID configurations, in addition to monitoring the status of RAID arrays and performing rebuilds.
- **Integrated Dell™ Remote Access Controller (iDRAC)**, a remote server-management processor embedded in 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. iDRAC9 helps manage PERC 12 cards. 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 the Dell Technologies command-line interface (CLI), known as RACADM. With iDRAC9, customers can also perform real-time storage operations.
- **Dell APEX™ AIOps** is an AI-driven observability and incident management software-as-a-service (SaaS) solution that simplifies IT operations, increases IT agility, and gives IT teams control over digital infrastructure. It helps simplify operations by transforming high-volume events and alert noise into actionable incidents with root causes and recommendations. It also enables collaboration across organizational silos to triage and resolve issues faster, and it helps improve control over service availability by predicting issues and intelligently automating remediation.

Summary and Conclusion

Hardware-based RAID consistently demonstrates its strengths in larger configurations and performance-critical environments. With its dedicated RAID controller, hardware RAID excels in delivering lower latency across all disk quantities, maintaining median latencies that are 99.9% lower than those of software RAID. This makes hardware RAID ideal for applications where consistent, low-latency performance is essential. Additionally, hardware RAID handles increasing I/O demands more efficiently as disk numbers rise, and it outperforms software RAID in terms of IOPS for configurations of four or more disks.

A key advantage of hardware RAID is its rebuild performance under load. Rebuilding a RAID array after a drive failure is crucial to maintaining data integrity. Our testing revealed that hardware RAID rebuilds maintained much higher rebuild performance under load: up to 26,243 times faster than software RAID in eight-disk configurations. This exceptional rebuild capability minimizes time vulnerable to additional disk failures while maintaining high performance during the rebuild process.

While hardware RAID provides stable, predictable performance for larger RAID setups, software-based RAID shines in smaller deployments. Software RAID capitalizes on system resources such as CPU and memory to achieve impressive scalability in bandwidth. It outperformed hardware RAID by as much as 3,521% in two-disk configurations.

Organizations deploying PowerEdge servers can capitalize on both approaches, maximizing performance with software-based RAID in smaller configurations while relying on hardware-based RAID for larger setups. The combination of Dell Technologies' server-management capabilities and support for AI workloads further enhances the potential for achieving optimal performance, making it easier to tailor solutions to specific operational needs.

Learn More

Learn more about Dell PowerEdge R760 servers at **https://infohub.delltechnologies.com**.
See how we conducted our tests in our **methodology report**.
See more research from **Prowess Consulting**.

Appendix: Hardware Configuration

	Server 1	Server 2	Server 3	Server 4
System	Dell™ PowerEdge™ R760	Dell™ PowerEdge™ R760	Dell™ PowerEdge™ R760	Dell™ PowerEdge™ R760
CPU	2 x Intel® Xeon® Platinum 8460Y+ processor (model 143, stepping 8)	2 x Intel® Xeon® Platinum 8460Y+ processor (model 143, stepping 6)	2 x Intel® Xeon® Platinum 8460Y+ processor	2 x Intel® Xeon® Platinum 8452Y processor
Total cores/threads per CPU	80/160	80/160	80/160	72/144
CPU frequency	2.0 GHz	2.0 GHz	2.0 GHz	2.0 GHz
Storage controller 01	Marvell Technology Group Ltd. Dell™ Boot-Optimized Server Storage (BOSS)-N1 monolithic	Marvell Technology Group Ltd. Dell™ Boot-Optimized Server Storage (BOSS)-N1 monolithic	Marvell Technology Group Ltd. Dell™ Boot-Optimized Server Storage (BOSS)-N1 monolithic	Marvell Technology Group Ltd. Dell™ Boot-Optimized Server Storage (BOSS)-N1 monolithic
Disk	1.6 TB Dell™ Enterprise NVMe® CM6 MU firmware 2.2.0	1.6 TB Dell™ Enterprise NVMe® CM6 MU firmware 2.2.0	960 GB Dell™ NVMe® PE8010 RI M.2 (SK hynix®)	480 GB Dell™ Enterprise NVMe® ISE 7400 RI M.2 (Micron Technology Inc.)
Number of disks	2	2	2	2
Storage controller 02	Multi-device RAID	Dell™ PERC S160 software RAID	Broadcom®/LSI® Dell™ PERC H965i Front (firmware: 8.4.0.0.18-29)	Broadcom®/LSI® Dell™ PERC H965i Front (firmware: 8.4.0.0.18-29)
Disk	1.6 TB Dell™ Enterprise NVMe® CM6 MU (KIOXIA Corporation) (firmware: 2.2.0)	1.6 TB Dell™ Enterprise NVMe® CM6 MU (KIOXIA Corporation) (firmware: 2.2.0)	1.6 TB Dell™ Enterprise NVMe® CM6 MU (KIOXIA Corporation) (firmware: 2.2.0)	1.6 TB Dell™ Enterprise NVMe® CM6 MU (KIOXIA Corporation) (firmware: 2.2.1)
Number of disks	8	8	8	8
Memory	8 x 32 GB (256 GB) DDR5, 4,800 megatransfers per second (MT/s)	8 x 32 GB (256 GB) DDR5, 4,800 MT/s	16 x 16 GB (256 GB) SK hynix® HMCG78MEBRA174, 4,800 MT/s single-rank	16 x 16 GB (256 GB) SK hynix® HMCG78MEBRA174N, 4,800 MT/s single-rank
Network	1 x 10/25/40/50/100 Gb/200 Gb Broadcom® NetXtreme® E-Series BCM57504 (rev 12)	Not applicable (N/A)—in Prowess Consulting's lab with no network card needed	2 x 100 Gb Broadcom® NetXtreme® E-Series P2100D BCM57508 QSFP PCIe®	N/A—in Prowess Consulting's lab with no network card needed
OS	Red Hat® Enterprise Linux® version 9.3	Windows Server® 2022 Datacenter Evaluation Desktop Experience version 21H2	Red Hat® Enterprise Linux® version 9.2	Windows Server® 2022 Datacenter Evaluation Desktop Experience version 21H2
BIOS version	1.5.6	1.5.6	1.5.6	2.1.5

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.

Prowess Consulting and the Prowess logo are trademarks of Prowess Consulting, LLC.

Copyright © 2025 Prowess Consulting, LLC. All rights reserved.

Other trademarks are the property of their respective owners.

0225/230210