The Micron® 9400 NVMe™ SSD Performance With NVIDIA® Magnum IO GPUDirect® Storage Platform

About This White Paper

This paper examines the performance and response time benefits of combining the Micron® 9400 NVMe™ SSD and NVIDIA® Magnum IO GPUDirect® Storage (GDS). It focuses on small (4KB), medium (128KB), and large (1024KB) transfer sizes in three different test scenarios, illustrating the Micron 9400 SSD (7.68TB)[1] advantages over the leading competitor's[2] performance-focused, 7.68TB NVMe SSD.

1. GDS compared to legacy IO performance: A legacy system (one without GDS) executes 4KB, 128KB, then 1024KB transfer sizes. Resultant performance is compared to the same workloads when GDS is in place.
2. An idle GDS system compared to a busy GDS system: This test is designed to illustrate the adverse performance impact when the test system is busy with an emulated artificial intelligence (AI) data pre-processing workload.
3. Micron 9400 SSD compared to leading competitor's SSD: This section uses a busy system for both comparisons. Results of this test show that the Micron 9400 SSD demonstrates higher performance and better response times compared to the leading competitor.

KEY FINDINGS

The Micron 9400 NVMe SSD brings performance and response time advantages to busy, GDS systems.

25% Higher Busy System Performance

The Micron 9400 SSD demonstrated up to 25% higher performance in a busy, GDS system when compared to the leading competitor (Competitor A).

Cuts Busy System Response Time in Half

The Micron 9400 SSD demonstrated latency up to 50% lower than Competitor A.

Figure 1: Micron 9400 NVMe SSD (U.2/U.3)

About the Micron 9400 NVMe SSD

The Micron 9400 NVMe SSD sets a new performance benchmark for PCIe® Gen4 storage. Decades of experience have led to an SSD that packs in up to 30TB of capacity. This SSD is designed for critical workloads like artificial intelligence training, high frequency trading, and database acceleration.

NVIDIA Magnum IO GPUDirect® Storage

NVIDIA Magnum IO GPUDirect Storage is a technology that enables a direct memory access (DMA) data transfer path between an SSD and graphics processing unit (GPU) memory as shown in Figure 2. NVIDIA is a pioneer in artificial intelligence, from introducing the GPU in 1999 to the first release of Magnum IO GPUDirect Storage (GDS) in 2021.[3]

Micron and NVIDIA have a long history of technical collaboration, including joint webinars focused on system architecture to overcome AI storage challenges.

1. Unformatted. 1 GB – 1 billion bytes. Formatted capacity is less.
2. Competitor identified in Forward Insights, SSD Insights Q4/22 analyst report.
3. Additional information is available here:https://www.nvidia.com/en-us/about-nvidia/corporate-timeline/ .

GDS Improves Performance up to 33% Compared to Legacy IO

GDS bypasses the CPU and DRAM, making direct memory access (DMA) data transfers between the SSD and GPU (Figure 2). The legacy IO path data transfer is more complex. The data moves from the CPU through a PCIe switch, to the CPU, then to the system memory (DRAM). From system DRAM, data moves back through the CPU, through the PCIe switch, then finally to the GPU (Figure 3).[4]

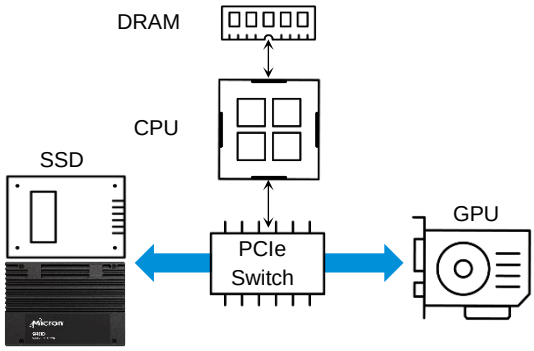

Figure 2: GDS IO Path

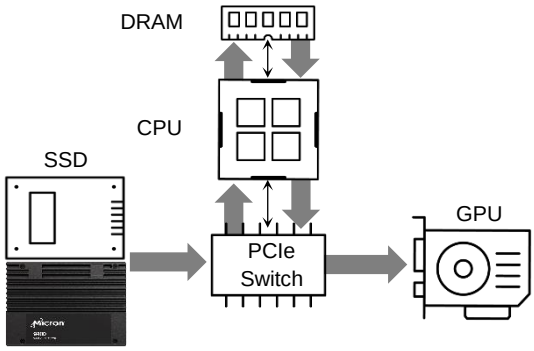

Figure 3: Legacy IO Path

Test data shows that the simpler, more direct data GDS path improves throughput and decreases response time. The greatest performance increase and response time decrease occur with small transfer sizes. Each dot on the lines in the following figures represents a specific queue depth at which performance and latency were measured (starting at a queue depth of 8 and doubling; e.g., 8, 16, 32, etc.). In the following performance and response time figures, comparison lines are drawn where performance difference is greatest.

Figures 4, 5 and 6 show the maximum GDS performance improvement and response time improvement for small (4KB), intermediate (128KB), and larger (1024KB) transfers. The specific queue depth with the highest variance is noted in the figures. Data points farther to the right on the horizontal axis indicate higher performance. Data points lower on the vertical axis indicate lower response time.

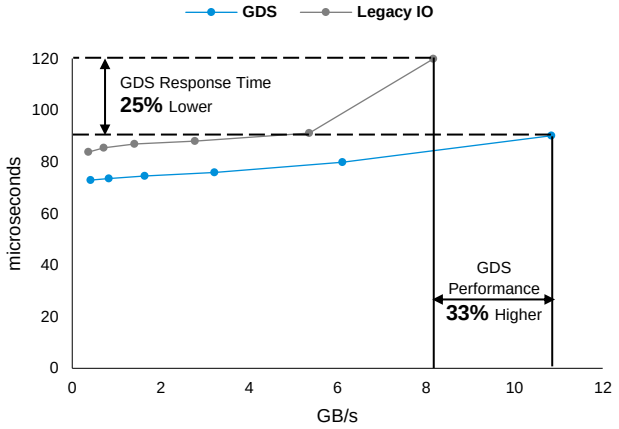

Figure 4: 4KB Transfer GDS Improvement (idle system)

4KB Transfers GDS Improvement: 33% Higher Performance and 25% Lower Response Time

GDS 4KB transfer size performance and responsiveness improvements are clearly visible in Figure 4. GDS increased 4KB transfer size performance by up to 33% when compared with legacy IO and improved 4KB transfer size response time by 25%.

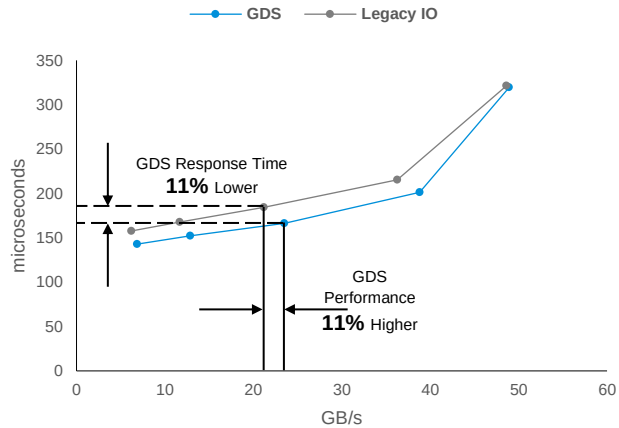

Figure 5: 128KB Transfer GDS Improvement (idle system)

128KB Transfer GDS Improvement: 11% Higher Performance and 11% Lower Response Time

GDS improved 128KB transfer size performance by up to 11% and improved 128KB transfer size response time by 11% as shown in Figure 5.

4. Additional information is available here: https://docs.nvidia.com/gpudirect-storage/overview-guide/index.html.

1024KB Transfer GDS Improvement: 16% Higher Performance and 16% Lower Response Time

GDS improved 1024KB transfer size performance by up to 16% and improved response time by 16% as shown in Figure 6.

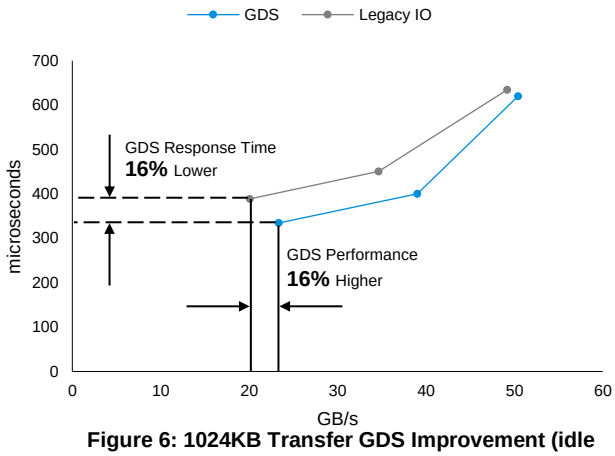

Figure 6: 1024KB Transfer GDS Improvement (idle system)

Adding an AI Data Pre-Processing Workload Reduces Performance, Increases Response Time

When complex, resource-intensive workloads like data pre-processing for AI model training run on a shared system, the data pre-processing workload can impact other workloads that run on the same system.

This section demonstrates the impact of an emulated data pre-processing workload on the 4KB, 128KB, and 1024KB workloads in a GDS-enabled system. The data pre-processing workload is generated using googlestresstestapp configured as noted in Table 1.[5]

System CPUs are highly loaded (86% utilization across both 64-core CPUs), and system memory is loaded at 205 GB/s (roughly 50% of the total available memory bandwidth) in this workload.

Loaded Subsystem	Loaded Value
CPU Utilization	86%
Memory Bandwidth Utilization	205 GB/s (about 50%)

Table 1: Emulated data pre-processing workload

Figure 7 shows the GDS data transfer path for the idle system (the system without the pre-processing workload). Figure 8 shows the GDS data transfer path and pre-processing workload execution location for the busy system (this inserted workload is what makes these systems ”busy”). These figures represent the configurations on which GDS idle and GDS busy system performance data is measured.

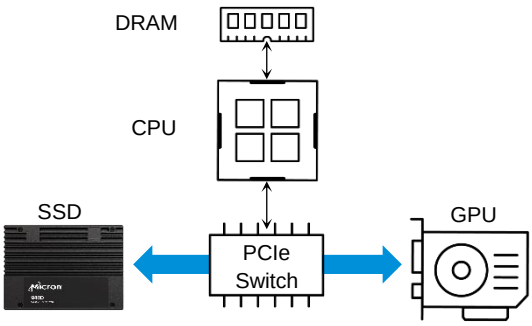

Figure 7: GDS IO Path, Idle System

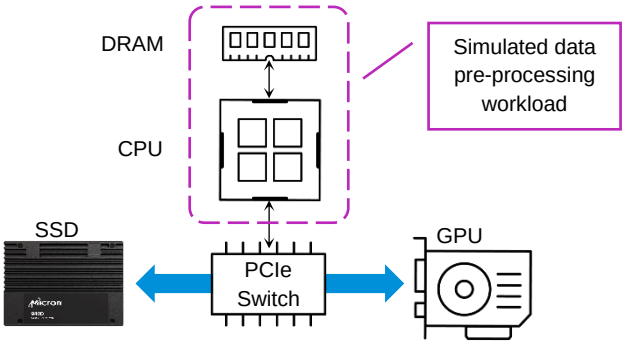

Figure 8: GDS IO Path and Simulated Pre-Processing Workload (Busy System)

5. Additional information is available here: https://github.com/stressapptest/stressapptest.

Figures 9, 10, and 11 show the performance and response time impact of adding the emulated data pre-processing workload. In each of these figures, introducing the data pre-processing workload decreases performance, increases response time, or both. This indicates that adding the data pre-processing workload shows a significant impact on all transfer sizes tested.

4KB Transfer: Busy GDS System Shows 50% Lower Performance and 2X Response Time

The impact of introducing the data pre-processing workload into a GDS-enabled system shows a significant impact on 4KB transfer size performance and responsiveness. Performance is halved (50% lower than the idle system), and response time is doubled relative to the GDS idle system as seen in Figure 9.

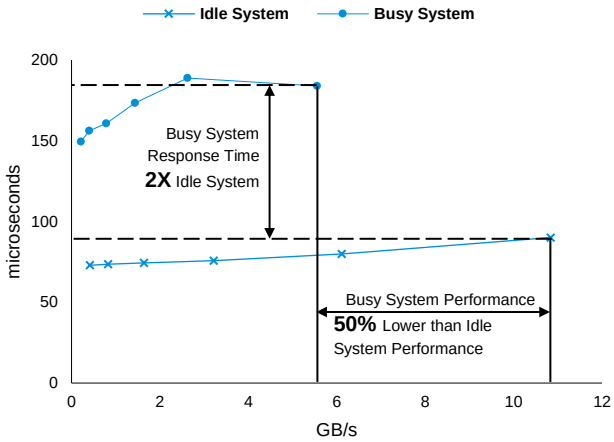

Figure 9: 4KB Transfer Idle vs. Busy GDS System

128KB Transfer: Busy GDS System Response Time is 34% Higher Than Idle System

Figure 10 shows that 128KB transfer size performance is nearly identical for the busy and idle systems, but the busy system response time is twice the idle system.

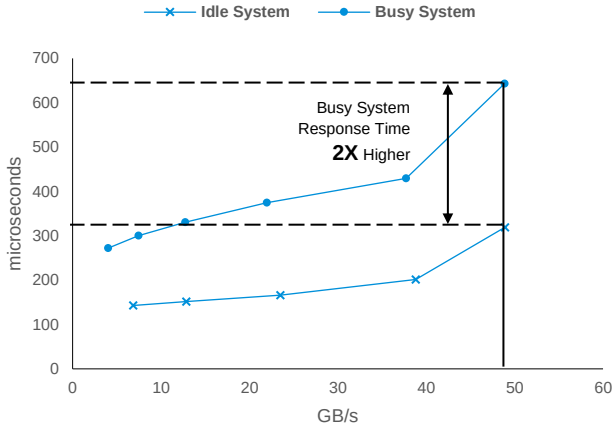

Figure 10: 128KB Transfer Idle vs. Busy GDS System

1024KB Transfer: Busy GDS System Shows 2X Response Time

Figure 11 shows that the busy GDS system response time is twice that of the idle GDS. The busy GDS system showed a negligible performance decline.

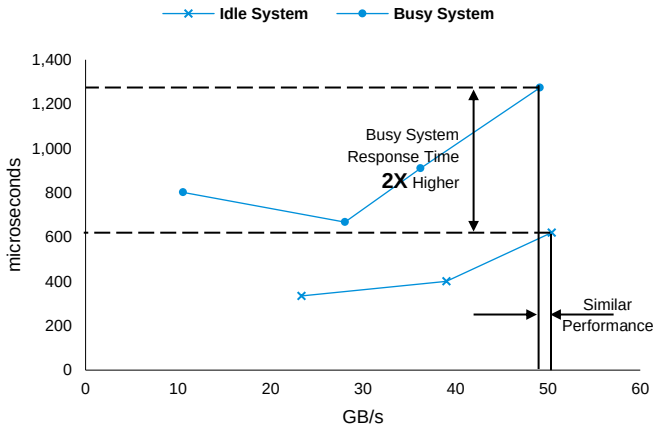

Figure 11: 1024KB Transfer Idle vs. Busy GDS System

Substantially Better Busy System GDS Results Than Competitor A

This section compares the GDS busy system performance and response time results of the Micron 9400 NVMe SSD and Competitor A. Figure 12 represents the busy GDS system with the Micron 9400 SSD while Figure 13 represents the busy GDS system with Competitor A.

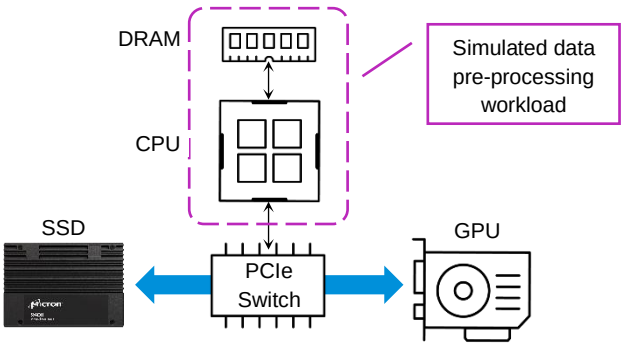

Figure 12: Micron 9400 SSD, GDS IO Path, Busy System

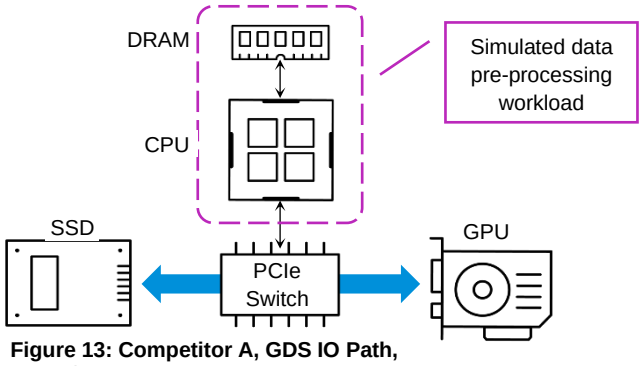

Figure 13: Competitor A, GDS IO Path, Busy System

4KB Transfer (Busy GDS System): Micron 9400 SSD Shows 25% Higher Performance and 23% Lower Response Time

The Micron 9400 SSD shows less impact with 4KB transfer sizes when the data pre-processing workload is introduced on the GDS-enabled system. The Micron 9400 SSD shows 25% higher performance and 23% lower response time than Competitor A as shown in Figure 14.

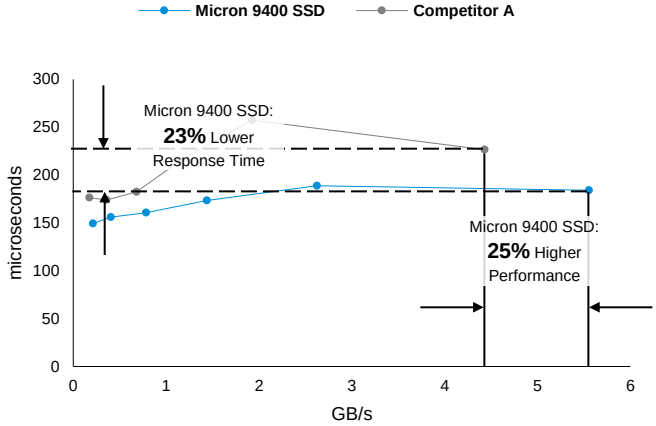

Figure 14: 4KB Transfer Busy GDS System Micron 9400 SSD vs. Competitor A

128KB Transfer (Busy GDS System): Micron 9400 SSD Shows 14% Higher Performance and 13% Lower Response Time

Figure 15 shows that the Micron 9400 SSD demonstrates 14% higher performance and 13% lower response time with a 128KB transfer.

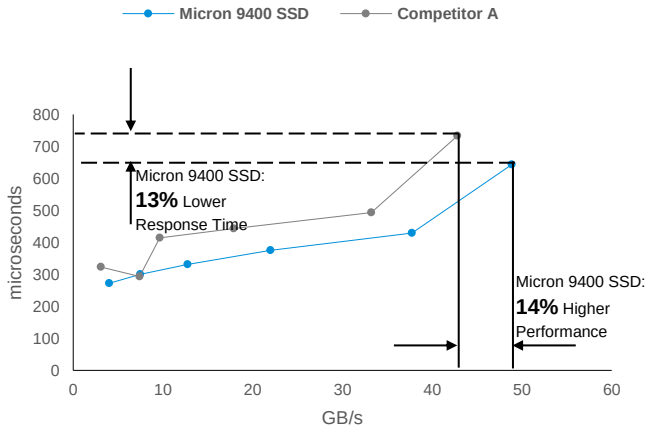

Figure 15: 128KB Transfer Busy GDS System Micron 9400 SSD vs. Competitor A

1024KB Transfer (Busy GDS System): Micron 9400 SSD Shows 2X Lower Response Time

Figure 16 shows performance and response time differences between the Micron 9400 SSD and Competitor A on a busy GDS system using a 1024KB transfer size.

Although the maximum performance of these two SSD is similar, Competitor A response time is double the Micron 9400 SSD at these maximum performance values.

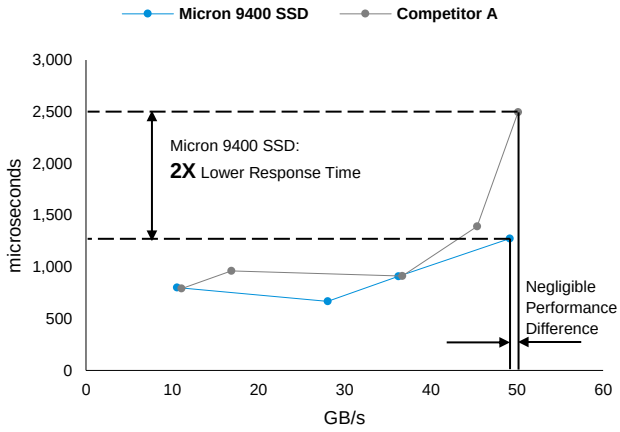

Figure 16: 1024KB Transfer Busy GDS System Micron 9400 SSD vs. Competitor A

Conclusion

NVIDIA Magnum IO GPUDirect Storage delivers performance and response time benefits for AI workloads,

The Micron 9400 sets a new performance benchmark for PCIe Gen4 storage, impacting real-world workloads. Micron is a market leader in the innovation and design of storage solutions that are built for today's advanced data centers. We are committed to helping customers maximize the potential in their data. Micron's decades of experience have led to an NVMe SSD that is optimized for the most demanding data center workloads, including those running Artificial Intelligence.

Combining these technologies shows that the Micron 9400 SSD enables superior GDS platform performance and response time over a competitor's performance-focused NVMe SSD. These results, while emulated, reveal that the Micron 9400 is the right choice for AI workloads that use data pre-processing.

micron.com/9400