Lab Insight

Western Digital OpenFlex™ Data24 – Shared High-Performance NVMe™ Storage

Author: Russ Fellows

October 2022

Evaluator Group

Western Digital

Overview

As cloud computing and other IT transformation projects continue to be implemented, companies are challenged with the need to attain the operational flexibility that cloud computing enables while still protecting their critical corporate data. Although clouds enable companies to dynamically align their costs with revenues, public clouds in particular present privacy and data governance challenges. As a result, IT consumers are seeking options that enable control and data governance while also providing operational flexibility.

One of the key technologies that enable operational flexibility is the ability to match IT resources with business requirements in real time. Storage requirements include delivering application performance and capacity needs while also meeting application availability and data protection requirements. Storage Area Networks (SAN) and other network accessible storage technologies became prevalent over 20 years ago due to the improved efficiencies of providing storage over shared access infrastructure.

Recent storage technology changes have driven some application owners to choose direct-attached storage rather than shared storage. While direct-attached solid-state storage can provide low latency, it lacks the high aggregate performance, efficiency or availability benefits of shared storage.

With this backdrop, Evaluator Group was asked to compare Western Digital's OpenFlex™ Data24, a disaggregated, external shared storage system to internal, server-based storage for application environments.

- Configuration A: Western Digital OpenFlex™ Data24 NVMe enclosure accessed as NVMe™ over fabric storage, with shared access to 6 servers
- Configuration B: Aggregate performance of direct-attached devices in 6 independent servers with 4 internal NVMe storage devices each

The two test environments were made as similar as possible, using the same servers, quantity and type of storage media for both configurations.

> ***Evaluator Group Comments: Testing was designed to fairly compare storage media located either in a server, or in an external OpenFlex™ Data24 enclosure. The Western Digital Data24 provides external NVMe storage that may either be directly attached to a single system or shared to enable greater storage utilization and flexibility, while also delivering performance approaching that of multiple systems with internal NVMe storage.***

Testing was performed using the synthetic "*fio*" (flexible i/o tool) workloads. The remainder of this paper explores the technology along with business and financial considerations of these two alternative storage deployments in more detail. A summary of results in Table 1 shows the performance of Western Digital Ultrastar® DC SN840 NVMe media located either in the OpenFlex Data24 shared enclosure or internally within six separate servers.

	OpenFlex Data24 IOPs	Internal-NVMe IOPs	OpenFlex Data24 %
4KB - 100% Write	6,262,935	6,261,094	100.0%
4KB - 70% Read	12,051,484	13,063,521	92.3%
4KB - 100% Read	15,270,557	18,478,808	82.6%
16KB - 100% Write	1,626,244	2,312,328	70.3%
16KB - 70% Read	3,042,418	4,185,171	72.7%
16KB - 100% Read	4,244,836	5,035,670	84.3%

Table 1: Western Digital OpenFlex Data24 vs. Internal I/O rate using Western Digital Ultrastar® DC SN840 media. (Source: Western Digital & Evaluator Group)

The Evolving IT Landscape

As outlined previously, several trends are driving significant changes within the IT landscape. The move towards utilizing public clouds and services is changing where IT operations occur and how they are paid for, as well as having an impact on how applications are managed. Additionally, new technologies that enable access to NVMe storage via a fabric provide significant availability, flexibility and performance benefits for companies that can leverage these options.

Clearly, one of the biggest changes to information technology has been the impact of the cloud, including both public and private cloud use. Cloud computing promises the ability to efficiently manage resource utilization and performance on-demand. By enabling applications to scale both up and down allows firms to optimize their efficiency. A key component of composable IT infrastructure is the ability to independently scale and connect server and storage resources.

The demand for ever improving business productivity and efficiency in turn drives the need for faster application processing while consuming fewer IT resources. As a result, businesses must improve their computational, networking and storage processing speeds or risk losing out to businesses that are able to drive business profitability through IT efficiencies. Storage is a key aspect of improving efficiency and flexibility, due in part to the number of choices of components and vendors offering solutions.

Storage Requirements

Efficiently managing IT resources requires companies to derive maximum benefits from an asset while minimizing the associated costs. Importantly, economics must be balanced against technology requirements to effectively meet application performance and availability requirements as economically as possible.

Storage Economics

There are several considerations that impact economics, including utilization, the service life of technologies along with practical aspects including internal asset management and accounting practices. While TCO models can provide insights, it is not necessary to construct complex models to understand that utilization of an asset during its lifespan has significant impact on the economic benefits delivered over its service period.

Another important consideration is the concept of the useful life of an asset, which in turn requires the ability to manage and track every asset. Historically, servers were often purchased along with internal storage. However, this is inefficient both technically and economically. First, servers and storage now have different useful lifespans and tying servers and storage together makes tracking these assets independently difficult.

Modern servers are typically replaced every three years, as this is the period when companies have decided that the benefits of new server technology are cost effective. For older storage systems comprised of spinning media, a three-year lifespan was also somewhat typical. However, with the widespread adoption of solid-state media storage, this is no longer true. Solid-state devices can now have a useful lifespan of 5 – 7 years, increasing the possible lifespan of storage systems to 5 years or more.

Finally, the practical consideration of asset tracking can have an impact on the choice of products. Technologies that are easy to use are more likely to be deployed than more difficult options. The time and complexity required to track expensive solid-state drives within a server makes using internal devices more complex. Although this aspect of economics is often overlooked, it is an important reality for companies who must manage IT resources that have become a significant portion of a company's assets.

All of these realities make the concept of adding solid-state devices to a server a difficult proposition from an asset tracking and resource management perspective. Although these tasks are straightforward to perform from a technical perspective, they are highly complex from an accounting viewpoint.

Storage Technology

Non-Volatile Memory Express (NVMe) is an interface that enables solid-state drives (SSDs) to communicate using the high-speed PCIe ® bus, which provides low-latency access to solid-state devices. NVMe supports billions of commands with significantly less delay or latency compared to SCSI or SATA devices. In order to accommodate communications to external devices, NVMe over Fabric extends the NVMe protocol to support accessing NVMe over a network fabric.

External storage provides both technical benefits along with better manageability and economics. As a result of the need to manage storage capacity, performance and storage assets separately from computing, many companies have a strong preference for utilizing external storage systems.

High-Speed NVMe Access Options

The ability to independently scale storage and computing capabilities is an important feature that enables greater flexibility, utilization and ultimately greater economic value. Although local storage can provide high speed access to data, it does so at the cost of reduced flexibility and utilization. Additionally, captive storage can become a stranded asset as it is literally tied to the server in which it resides. In summary, a table comparing internal NVMe devices to network accessed, shared NVMe devices is presented below.

	Western Digital Data24	Internal NVMe	Advantage
Full Speed Access to NVMe media	Full speed for small block I/O	Full speed for small and large block I/O	Internal – Better for large block I/O
Management API for Provisioning	Yes	No	Data24 – Improved manageability
Independently Scale Compute & Storage	Yes	No	Data24 - Increased flexibility and utilization
Multiple Data Paths	Yes	No	Data24 - High availability

Table 2: Western Digital Data24 vs. Internal NVMe media. (Source: Evaluator Group)

Western Digital Data24 - Overview

Western Digital's OpenFlex Data24 NVMe-oF ™ storage platform provides the high-performance of server resident NVMe flash storage, along with the benefits of external, shared storage. The Data24 system provides low-latency sharing of NVMe SSDs over a high-performance Ethernet fabric to deliver similar performance to locally attached NVMe SSDs. Utilizing Western Digital Fabric adapters, the Data24 allows up to six hosts to be attached without a switch or up to 48 hosts when sharing storage through one or more switches.

The OpenFlex Data24 system can be used as a dedicated or shared storage resource in composable infrastructures, providing high availability and enterprise-class storage reliability. Applications often deployed include real-time analytics, databases, virtualized server applications, AI/ML workloads and high-performance computing.

Data24 features include:

- High density and capacity, up to 368TB* in 2U
- Low latency storage access (as low as 100 ms)
- Dual Port SSDs
- Bandwidth match between SSDs (Storage) and I/O (Network) – No oversubscription
- RESTful management API
- Global service and support
- Limited warranty of 5 years

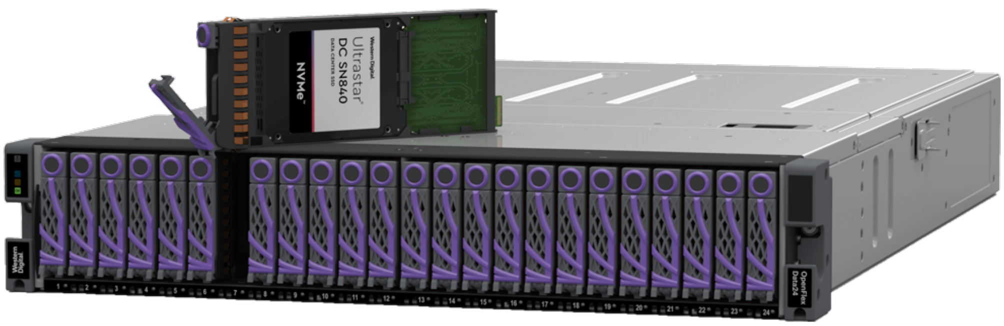

Figure 1: Western Digital OpenFlex™ Data24 (Source: Western Digital)

Western Digital Data24 Performance Testing

Evaluator Group was asked to analyze the performance and features of two different storage environments able to support NVMe access to applications. Working with Western Digital, we setup two comparable storage environments in the Western Digital labs. All testing was performed by Western Digital working with Evaluator Group to review the testing process, along with providing Evaluator Group full results to all tests. All hardware, software and other aspects are documented in the Appendix.

Storage performance was measured using the ‘fio’ tool to create workloads. Testing consisted of 8 different workloads, utilizing three different block sizes and multiple read – vs. write ratios:

- 4KB, 0%, 70% and 100% read, with 100% random access
- 16KB, 0%, 70% and 100% read, with 100% random access
- 128KB, 0% and 100% read with 100% sequential access

* *One terabyte (TB) is equal to one trillion bytes. Actual user capacity may be less due to operating environment*

These workloads were used to compare the performance of a Western Digital Data24 enclosure, accessed via a NVMe over Fabric on 100 Gb/s Ethernet connections to six servers, compared to those same six servers accessing internal NVMe storage devices. Both test environments utilized the exact same number, make, model and capacity of NVMe storage media. The testing used 24, Western Digital Ultrastar DC SN840-3.2 TB NVMe SSDs. A diagram of the test environment, along with additional details on the tests can be found in the Appendix.

Evaluator Group Comments: Western Digital has developed rigorous test procedures that yield reproducible results that are highly consistent across multiple test runs. Evaluator Group was able to verify that the results obtained during the audited testing are very similar to past testing performed by Western Digital. As a result, we can confidently state that customers should be able to achieve these results in their environment.

For specific details regarding the configuration of the Data24 vs. internal storage configurations, please refer to the Appendix.

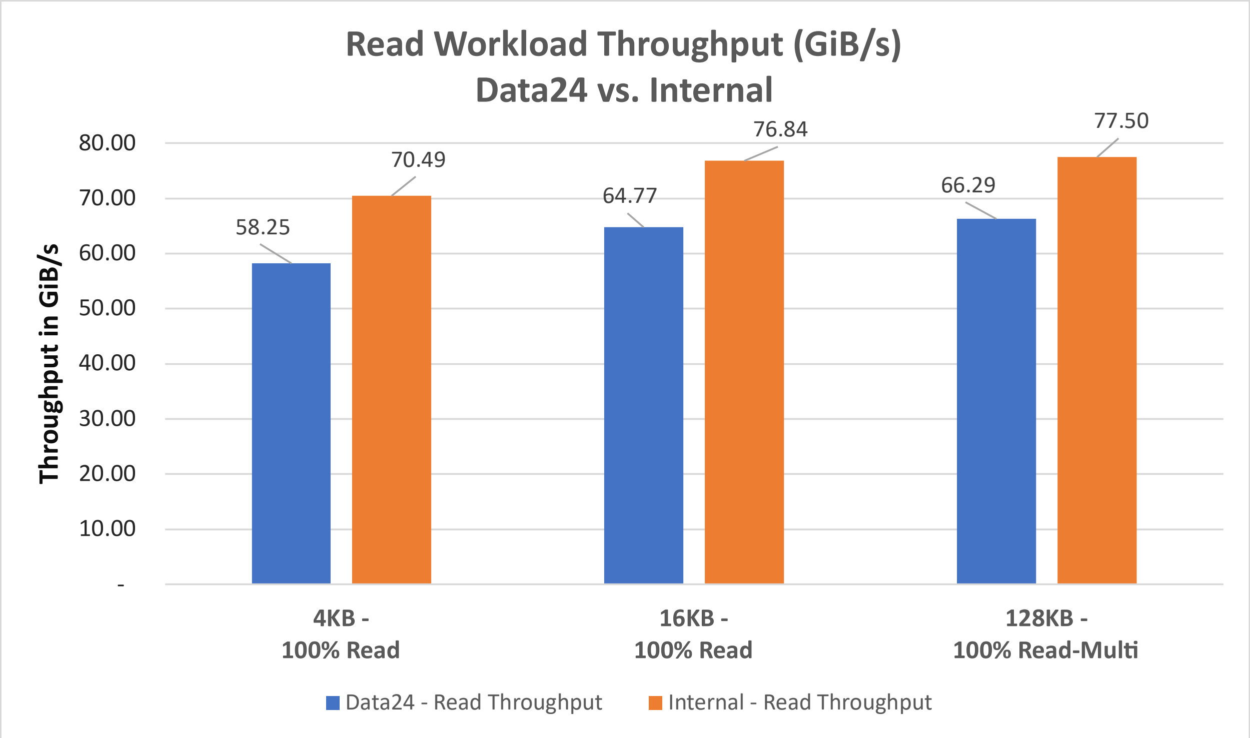

Figure 2: Data24 vs. Internal NVMe – Read Throughput Comparison (source: Western Digital & Evaluator Group)

Figure 2 above shows the throughput results for 3 different block sizes for 100% read workloads for both test environments, using the exact same quantity and model of NVMe media. For the 128KB workload, multiple processes were utilized to increase the overall throughput, hence the title "Read-Multi."

The throughput results of different block sizes may be directly compared. The ability to achieve nearly the same throughput for small block access highlights the efficiency and low latency of both storage environments.

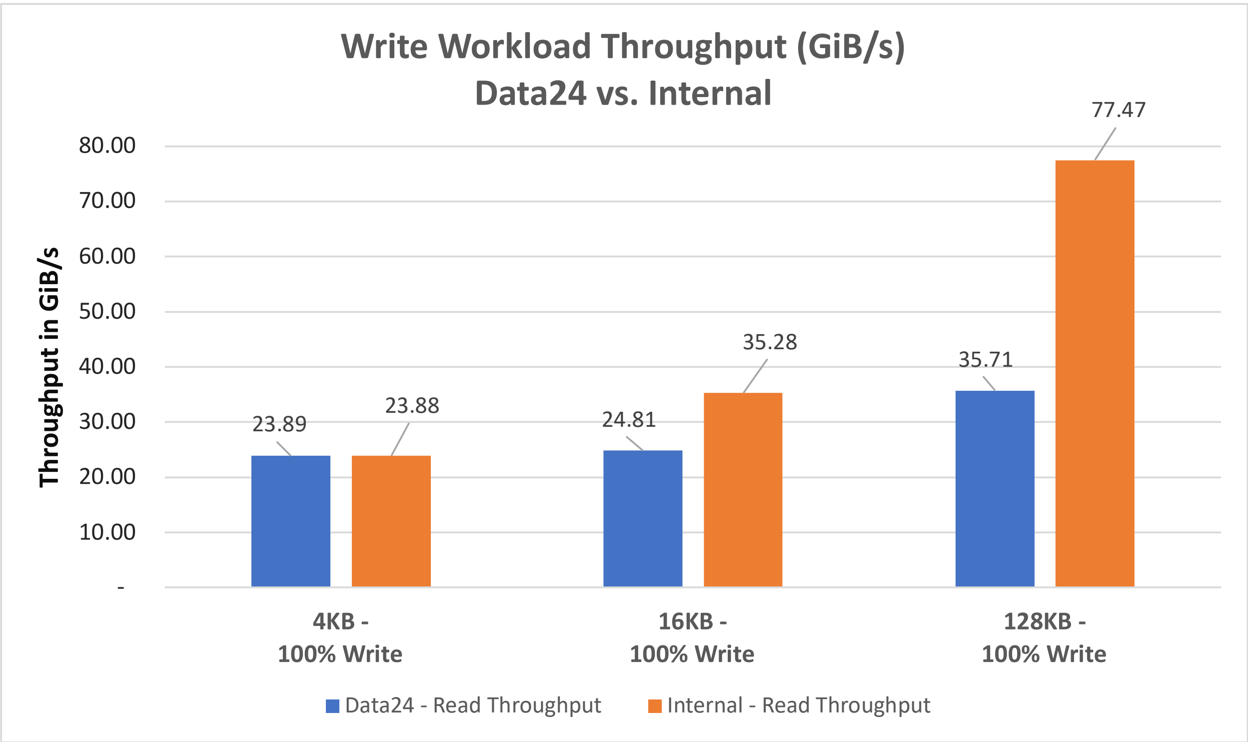

Figure 3: Data24 vs. Internal NVMe – Write Throughput Comparison (source: Western Digital & Evaluator Group)

Similarly, Figure 3 shows the throughput results for 3 different block sizes for 100% write workloads for both environments, using the same quantity and model of NVMe media. For smaller block sizes, Data24 has similar performance, although for large block sequential access the internal scenario was able to deliver slightly more than 2x the performance of the Data24.

Evaluator Group Comments: Western Digital Data24 system delivers low latency and high I/O rate performance that exceeds many enterprise class storage arrays, without the added cost or complexity of array-based data protection features. The throughput levels for each of the workloads tested exceeds the performance of all other storage systems tested to date by Evaluator Group.

Performance Details

The I/O rate (measured as I/Os per second, or IOPs) is a useful measurement for workloads up to 32 KB, although comparisons between different block sizes is more complex. Shown in the next three figures are comparisons of 4 KB, 16 KB and 128 KB workloads. Each figure shows a single block size with the resulting I/O rates and latency plotted for both internal and Data24 test environments for all workload mixes tested.

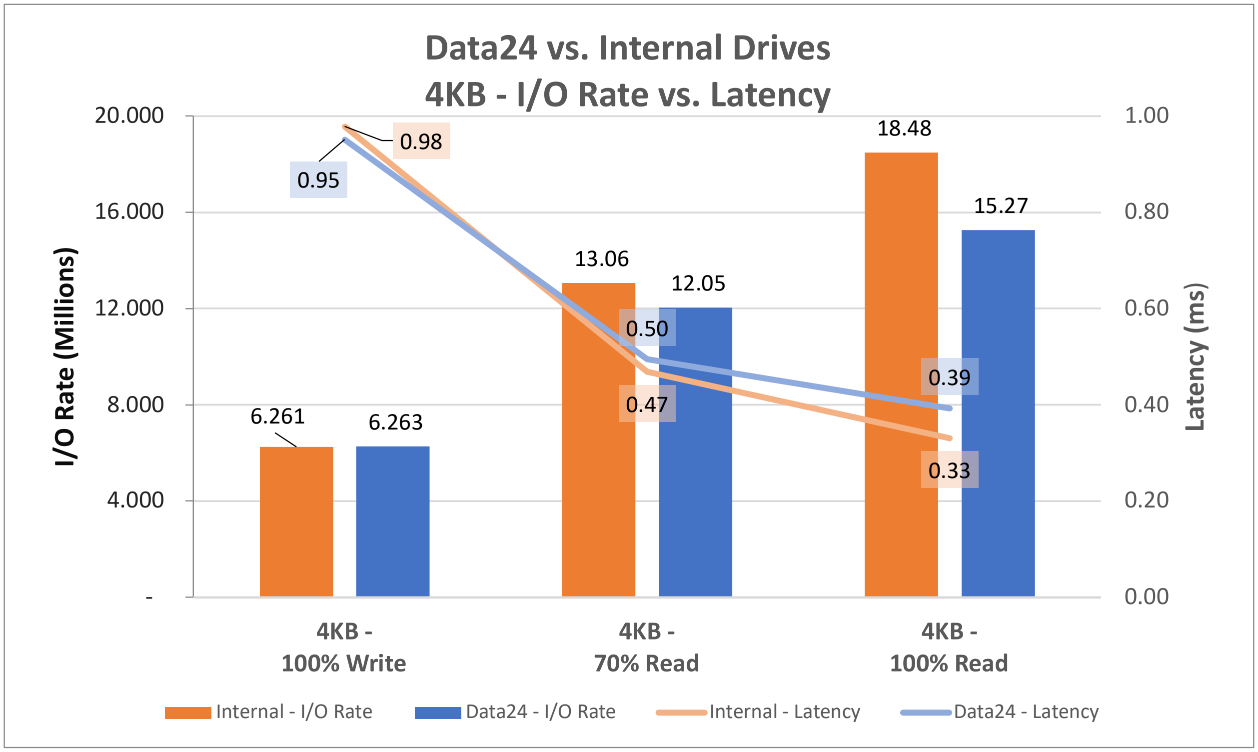

Figure 4: Data24 vs. Internal NVMe – 4 KB I/O Rate Comparison (source: Western Digital & Evaluator Group)

The three workloads shown above are very common with databases or other transactional applications. Although many databases attempt to perform 16K or 32K I/O operations, in many cases they will use smaller sizes rather than waiting to coalesce small I/Os. Additionally, databases often perform 4K transactions for highly latency sensitive I/O such as transaction logging operations.

Shown below are I/O rate comparisons for 16KB 100% read, 70% read and 0% read operations for both Data24 and Internal NVMe access.

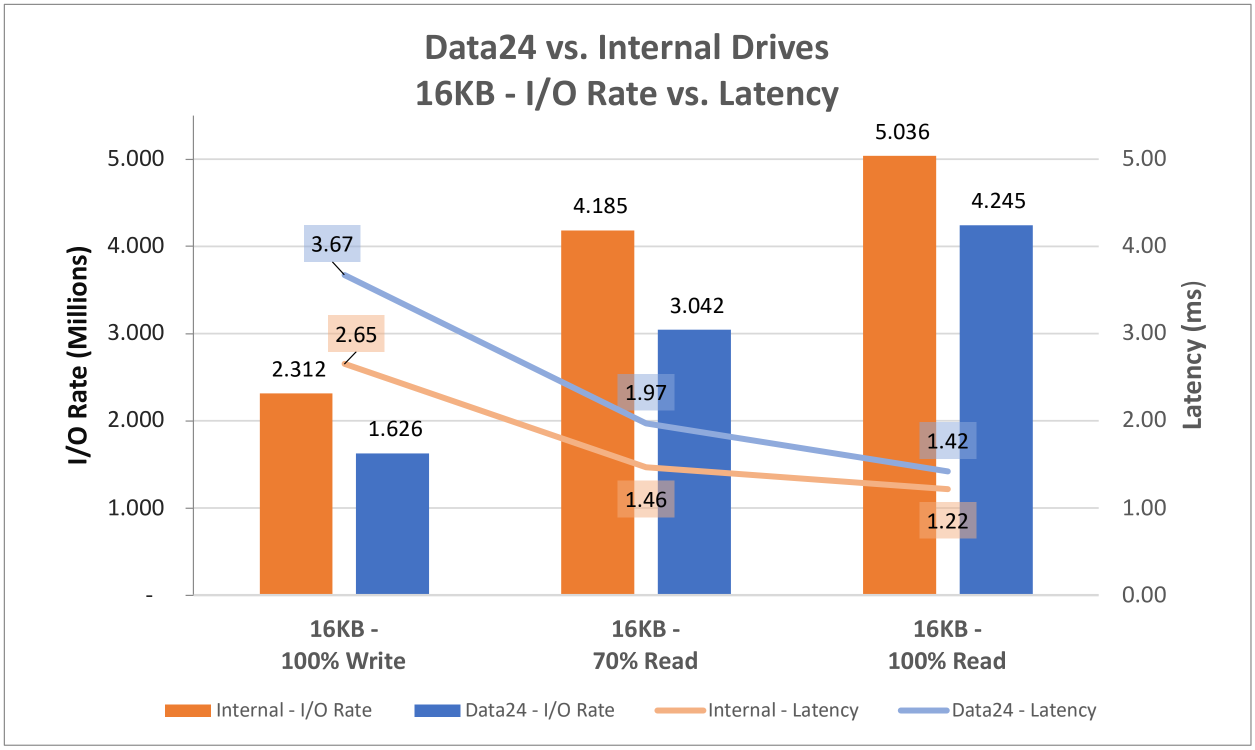

Figure 5: Data24 vs. Internal NVMe – 16 KB I/O Rate Comparison (source: Western Digital & Evaluator Group)

No approximation is perfect for predicting real world applications, however; many databases are designed to perform a significant portion of I/O at 16 or 32 KB. The read vs. write ratio will depend upon the application workload, and thus the three different ratios are shown.

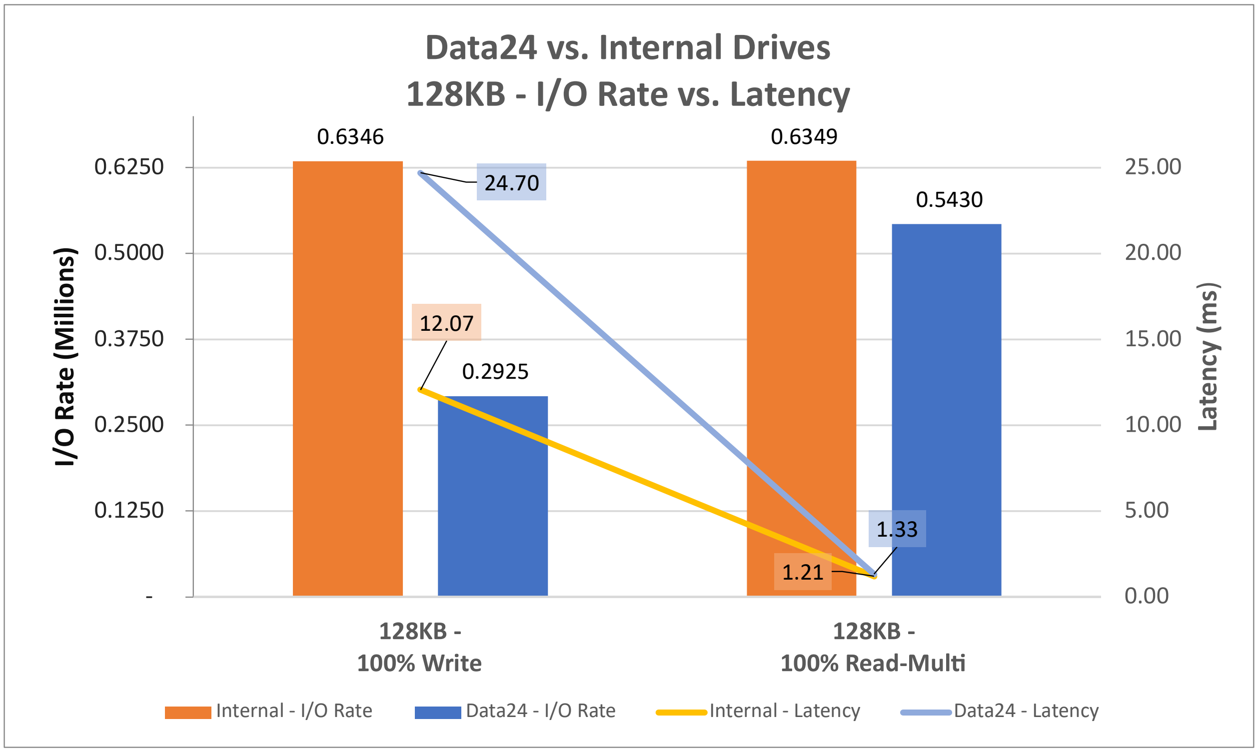

Figure 6: Data24 vs. Internal NVMe – 128 KB I/O Rate Comparison (source: Western Digital & Evaluator Group)

In addition to smaller block workloads, some applications utilize large block sizes for I/O operations, at sizes of 128 KB or larger. Although there can be some modest efficiency gains from using blocks larger than 128 KB, they are often very minor. Moreover, 128 KB is widely viewed as being representative of a systems maximum effective throughput rate for applications.

As Figure 6 indicates, the Data24 delivered more than 85 % of the performance of internal NVMe devices for a read workload. However, the write performance was roughly half that of the internal NVMe configuration. The large block write workload was the only comparison where a significant differential was observed.

Evaluator Group Comments: Across multiple workloads, the Western Digital OpenFlex Data24 system performed well, nearly matching the aggregate performance of 24 direct-attached NVMe devices, distributed across multiple servers. For latency sensitive applications that don't require array-based data protection features, the Western Digital Data24 is a highly compelling shared storage solution.

Final Thoughts

There are many important considerations for enterprise storage, including cost, performance, availability, reliability and data protection features. A growing number of applications no longer rely upon hardware to deliver resiliency or data protection. As a result, it is possible to deliver optimal levels of storage performance together with the flexibility of shared storage using pooled storage media. By eliminating unnecessary storage array features, while still providing shared storage access, companies can meet their applications storage requirements cost effectively.

Although data availability features may not be required, the ability to access storage from multiple servers provides flexibility to disaggregate servers and storage. Applications such as transactional databases, data warehouses along with many containerized micro-service applications utilizing Kubernetes are designed to provide their own data availability and redundancy without relying upon the underlying hardware.

Evaluator Group Comments: For applications that do not require array-based data protection, shared access to NVMe devices can provide storage pooling and independent scaling without the cost or performance impact of storage arrays. Western Digital's Data24 provides shared storage access and performance that surpasses traditional arrays without unnecessary availability or protection features.

Organizations may choose to utilize external storage arrays with enterprise availability and data protection services for specific applications requiring special protection. However, for many other applications, shared access storage enclosures provide an option for companies seeking the most cost-effective option for delivering high-performance storage while still enabling servers and storage to scale independently.

When storage software or applications provide the appropriate levels of data availability, additional storage features from a storage array add cost, while reducing performance. Western Digital's Data24 provides shared storage access and performance that surpasses traditional arrays without unnecessary availability or protection features.

Appendix

Test Environment Details

The test environment utilized the following hardware software and application workload items. Testing was performed in Western Digital's labs in July 2022.

Hardware and Infrastructure

Figure 6 below depicts the two test environments

- Common elements for both test environments
  - 6 Lenovo ® SR650 Servers, each with 2x Xeon ® Gold 6154 CPU and 384 GiB DRAM
  - External NVMe – with Western Digital OpenFlex Data24
- Western Digital OpenFlex Data24 connected via 6 x 100 GbE
  - 24 Ultrastar DC SN840 NVMe drives (3.2TB, 3 DW/D)
- Internal NVMe – connected to NVMe via internal PCIe
  - Total of 24 Ultrastar DC SN840 NVMe drives (4 per server)

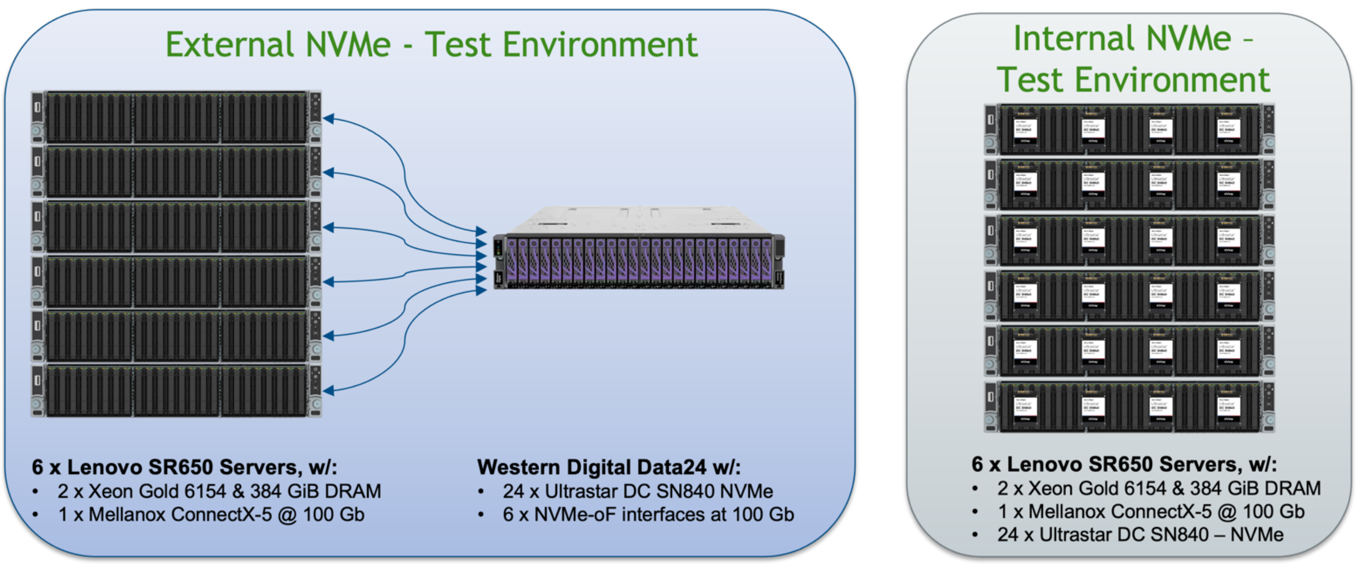

Figure 7: Test Environments for Western Digital Data24 and Internal NVMe devices (source: Evaluator Group)

Software Environment

- OS version: CentOS 8.4.2105
- NVMe version: 1.16
- MOFED driver: 5.6.1

Storage Configuration "External NVMe-oF"

- Six Lenovo SR650 systems connected to internal NVMe drives
- Total of 24, Western Digital Ultrastar DC SN840 NVMe SSDs
  - Note: 4 NVMe SSDs per system as shown in Figure 7

Storage Configuration "Internal NVMe"

- Six Lenovo SR650 systems connected to a single Western Digital OpenFlex Data24 system
- Total of 24, Western Digital Ultrastar DC SN840 NVMe SSDs
  - Note: All 24 NVMe SSDs within Western Digital OpenFlex Data24 enclosure as shown in Figure 7

Application Workloads

- The "fio" tool was used to generate synthetic workloads with different block sizes and read to write ratios
- Each test phase consisted of the following steps:
  - Each "device" is formatted
  - Next a "write-fill" workload to completely write the entire storage capacity tested
  - Each workload is then run for a 10 minute "warmup" period
  - Finally, a 20 minute "measurement" period is run
- Each test is run sequentially, with the completion of all 5 tests called a "set"
- Each test set was run 3 times, with the average of the 3 runs used for comparison
- The 8 workloads previously outlined were run using both storage configurations:
  - 4-KB, 100% read, with 100% random access
  - 4-KB, 70% read / 30% write, with 100% random access
  - 4-KB, 100% write, with 100% random access
  - 16-KB, 100% read with 100% random access
  - 16-KB, 70% read / 30% write with 100% random access
  - 16-KB, 100% write with 100% random access
  - 128-KB, multi-process 100% read, with 100% sequential access
  - 128-KB, 100% write, with 100% sequential access

Overview of Test Process

- Install Linux on all 6 servers
- Configure NVMe devices with 2 namespaces each, for a total of 48 namespaces across 24 devices
- Configure each server with 8 NVMe namespaces, residing on 4 NVMe devices
- Run fio workloads on each workload server, distributing workload to all 48 NVMe namespaces
- Gather the result sets for each workload to find average I/O rate, throughput and latency

About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures, and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end-users and vendor professionals through product and market evaluations, competitive analysis, and **education**. **www.evaluatorgroup.com** Follow us on Twitter @evaluator_group*

This document was developed with Western Digital funding. Although the document may utilize publicly available material from various vendors, including Intel and others, it does not necessarily reflect such vendors' positions on the issues addressed in this document.