lenovo

Analyzing the Performance of Lenovo DDR3 eXFlash DIMMs

Introduces the memory-channel storage offering from Lenovo

Describes the technical features of the devices

Explains the testing methodology used to measure the eXFlash DIMMs

Analyzes the results from evaluations performed

Tristian “Truth” Brown

LENOVO PRESS

Abstract

Lenovo eXFlash DIMMs are a memory-channel storage solution for the System x and Flex System servers. They are high performance, solid-state devices that are installed in regular memory DIMM slots and use the standard low-profile DIMM form factor. The result is a highly scalable flash storage device with low latency and high bandwidth capabilities.

eXFlash DIMMs are currently available in 200 GB and 400 GB per DIMM capacities. Up to 12.8 TB of total capacity can be installed in supported systems.

This paper provides a brief overview of the technology that supports eXFlash memory-channel storage, and quantifies the performance capabilities of the product. The paper focuses on the low-level hardware performance of single and scaled, 400 GB eXFlash storage devices.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, which provides information and best practices for using Lenovo products and solutions to solve IT challenges.

For more information about our most recent publications, see this website:

http://lenovopress.com

Contents

Introduction

In today's hyper-competitive business world, high-speed data transmissions and transactions are paramount in gaining a tactical edge over the competition. Over time, we saw significant improvements in processor performance with the introduction of increased frequencies and increased core-count.

The cost per gigabyte of storage also rapidly diminished, and some day might no longer be measured in dollars, but in cents. With these improvements in processor performance and storage costs, we see improved overall performance capabilities. However, newer customer applications are challenging us to keep pace in the areas of cloud analytics, big data, and high-frequency financial transactions. These applications require *balanced servers* that can keep pace with the speed of these high-volume, high-output processors.

The most efficient way to improve server balance is to focus on low-cost bottlenecks of the servers. In many cases, these bottlenecks are related to storage I/O activity because the increased speeds in storage technologies are still outpaced by the processing capabilities of servers. As a result, multi-core processors must wait for information that was requested from storage.

This relationship can be observed in a database system. For some database systems, there is a delicate balance between CPU processing power and I/O throughput from storage. When CPU processing power is added without a corresponding increase in storage I/O performance, this addition constrains the performance of the database system because of the wait time by the CPU for I/O from the storage device.

As storage I/O performance is improved, the configuration is better balanced and is more likely to meet the I/O demands of database systems. Figure 1 shows the delicate balance between storage IO per second (IOPS) and CPU processing capacity.

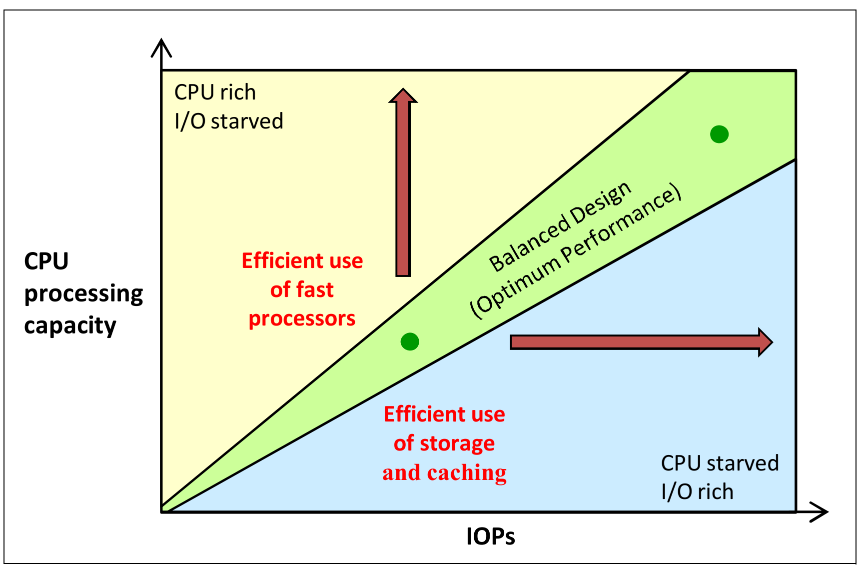

Figure 1 Trade off between server processing capacity and storage I/O

Internal server storage can be classified in the following form factors:

- Traditional spinning drives
- Solid-state memory

Traditional spinning drives are made up of hard disk drives (HDDs). Solid-state memory devices come in many form factors, such as solid-state drives (SSDs), PCIe flash adapters, or flash DIMMs. Spinning drives are appealing because of the massive amounts of data that can be stored at a relatively low cost per GB.

SSDs have a higher initial cost, but require a smaller overall footprint for equivalent data capacity. SSDs require the use of redundant array of independent disk (RAID) controllers and disk enclosures to function.

Flash DIMM technology that is implemented on the memory bus does not require more hardware that is associated with other industry-standard flash technologies, such as external or internal enclosures, platform controllers, SAS/SATA controllers, or RAID controllers. It is because of this reason that eXFlash memory-channel storage solution shines.

Figure 2 shows a simplified view of the inherent performance difference between equal capacity memory-channel storage devices, SSDs, and HDDs.

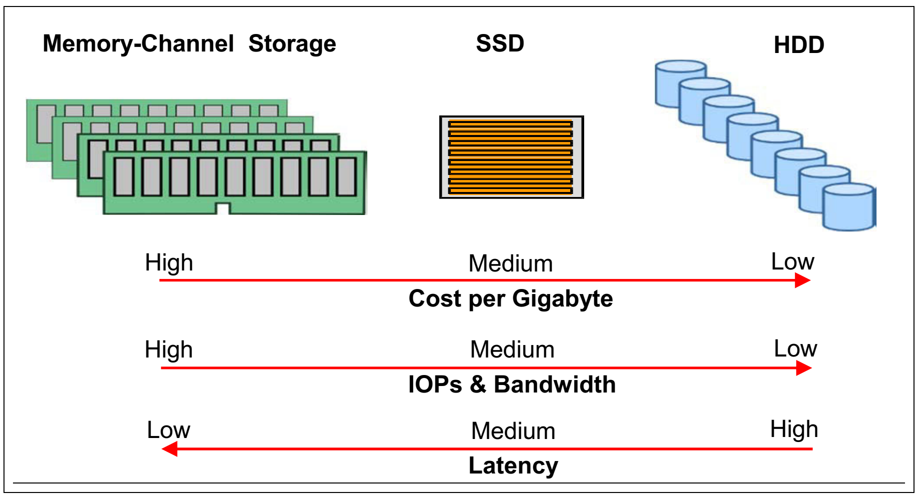

Figure 2 Performance differences between equal capacity storage technologies

Typical server workloads can be defined as a single or combination of the following types:

- IOPS-intensive applications require that storage devices to quickly process many read and write operations, with varying data transfer sizes. The capability of a device is quantified by measuring IOPS at transaction block sizes of 4 K, 8 K, and 16 K. Instances of these workloads are most common in public clouds, online transaction processing (OLTP) databases, and virtualized applications.
- Bandwidth-intensive applications require that storage devices transfer significant amounts of data by using large block sizes. Common transaction block sizes are 64 K - 1024 K. Transfer rates for individual devices are measured in megabytes per second (MB/s), while larger aggregate systems are measured in gigabytes per seconds (GB/s). Media streaming, file servers, and data backup activities are typical applications that drive these types of high throughput workloads.

- Latency-sensitive applications are applications that operate on transactions with smaller block sizes and that tend to exercise lower queue depths to minimize latency between transactions. Latency is the amount of time an application waits for an I/O operation to be completed by a storage device. For solid-state technology, this time often is measured within the mid to low micro-second (µs) range. Big data analytics, web serving, and high-speed financial transactions are examples of workloads that are sensitive to latency between transactions.

Lenovo's eXFlash memory-channel storage is an industry-first, solid-state storage technology that uses a standard DIMM form factor to significantly increase server and storage performance. It also efficiently helps to eliminate storage I/O bottlenecks for the workloads that are listed in Table 1.

Table 1 Typical application workload patterns

Application type	Workload type							
	IOPS intensive	**Bandwidth intensive**	**Latency sensitive**	**Read intensive**	**Write intensive**	**Random access**	**Sequential access**	**Good for SSD**
OLTP database	Yes	No	Yes	Yes	Yes	Yes	No	Yes
Data warehouse	No	Yes	Yes	Yes	No	Yes	No	Yes
File server	No	Yes	No	Yes	No	Yes	No	No
Email server	Yes	No	Yes	Yes	Yes	Yes	No	Yes
Medical imaging	No	Yes	Yes	Yes	No	Yes	No	Yes
Video on demand	No	Yes	Yes	Yes	No	Yes	No	Yes
Web/Internet	Yes	No	Yes	Yes	No	Yes	No	Yes
Web 2.0	Yes	No	Yes	Yes	Yes	Yes	No	Yes
Archives/backup	No	Yes	No	No	Yes	No	Yes	No

eXFlash DIMM features

Lenovo's eXFlash DIMM technology uses DDR3 memory interface technology with a compact form factor to provide a dense, high-performance storage solution. eXFlash DIMMs can help reduce total cost of ownership while improving system response time and I/O capabilities.

The following examples show how eXFlash DIMMs can enhance server performance:

- Maximizes storage footprint with unused DDR3 memory slots:
  - Expand available storage capacity on a server without adding traditional internal or external storage devices
  - Flash is accessed through the industry-standard DDR3 interface
  - Interoperable with standard RDIMMs in the same DDR3 channel
  - Uses software drivers for major x86 operating systems
- Provides high I/O performance with near linear scalability:
  - Drives multiple eXFlash DIMMs without experiencing performance degradation
  - Provides high-bandwidth performance that is independent of transaction block sizes
  - Provides excellent burst performance across mixed workloads

- Offers ultra-low write latency with Lenovo WriteNow technology (see Figure 11 on page 15):
  - Reduces hardware response time between workload transactions
  - As low as 5 µs response time for write operations

For a specific workload and queue depth, the performance of a eXFlash can easily be determined. This factor is important when you are extrapolating performance that is based on application needs.

Product description

eXFlash memory channel storage is a flash storage technology offering that is available for selected models of the System x® and Flex System™ family of servers. eXFlash DIMM is a high-performance, solid-state device that operates on memory DIMM slots by using the standard LPDIMM form factor. eXFlash DIMMs are fundamentally closer to the processor data stream and they do not require onboard controllers or PCIe interfaces for data transmission. This technology enables the eXFlash DIMMs to access information at much faster rates with lower latencies.

Figure 3 shows the conceptual implementation differences between industry-standard flash storage technologies.

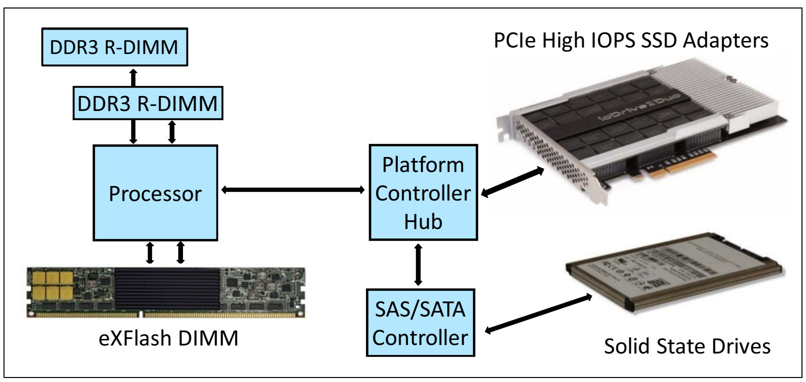

Figure 3 Implementation differences between industry storage technologies

eXFlash DIMMs address key data center infrastructure growth areas by realizing the following benefits:

- Decreasing write latency
- Providing ultra-fast temporary file storage
- Improving I/O read caching (for example, in VMware and KVM)
- Extending memory IO buffering (ultra-fast paging)

Lenovo's eXFlash DIMMs can help to achieve the following goals:

- The highest IOPS density per GB and the potential for the lowest write latency among Lenovo Flash storage offerings
- Linear storage I/O performance scalability for read/write intensive enterprise applications
- The ability to virtualize data-intensive enterprise applications that are limited by traditional storage I/O constraints
- Higher reliability and availability of services because of the fewer number of components that are exercised in the implementation of the solution

Technical specifications

The eXFlash DIMM is visually similar to a standard RDIMM, with the most notable differences being the heat sink and the onboard circuitry for the flash memory implementation. This section examines some of the factors that give the eXFlash DIMM the unique ability to function on a DDR3 memory channel.

Figure 4 shows the Lenovo eXFlash DIMM product.

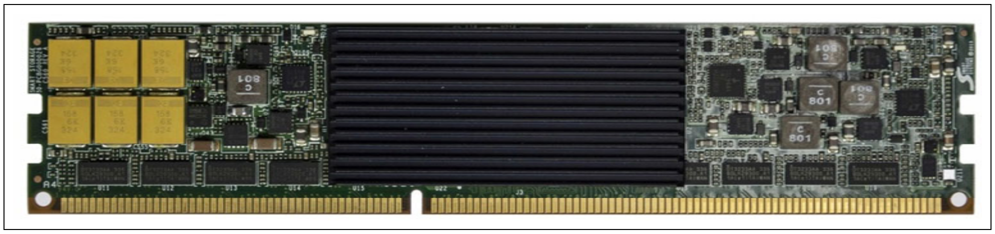

Figure 4 Lenovo eXFlash DIMM

Lenovo eXFlash DIMMs are recognized by the server as block storage, similar to other industry solid-state storage devices. However, a unique kernel driver is required for the operating system to use eXFlash DIMMs. The unified extensible firmware interface (UEFI) on the server logically isolates the traditional DRAM and eXFlash DIMM memory address spaces. A small amount of system memory is reserved for eXFlash DIMM operations.

As applications send storage I/O requests to the operating system, the data is forwarded to the eXFlash DIMM kernel driver. The kernel driver then performs the respective memory-channel commands to access the data that is stored on the eXFlash DIMM. Next, the requested data is directly transferred from flash memory to system memory by using a DDR3 memory bus.

Figure 5 shows a block diagram of this process.

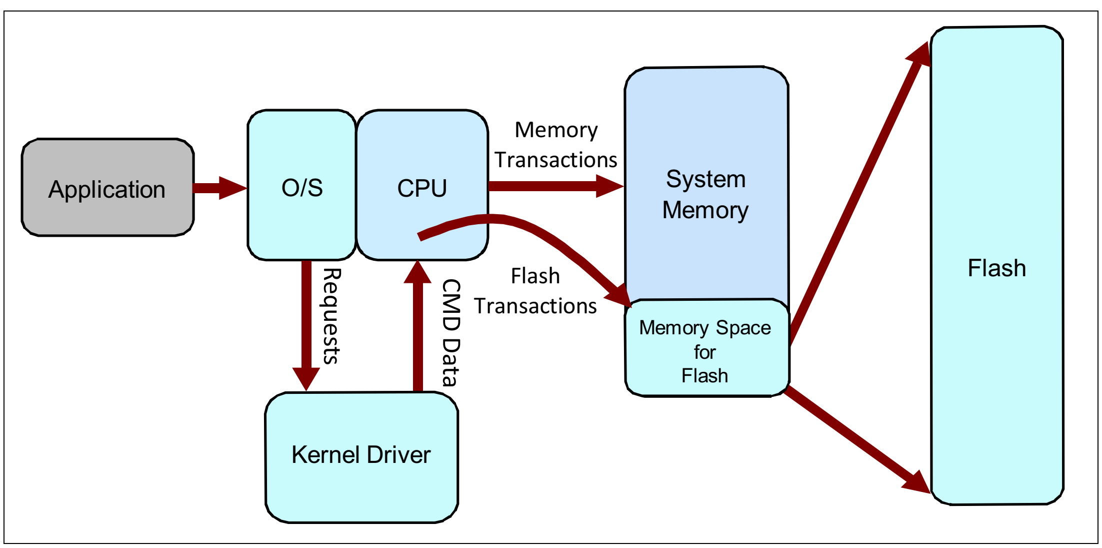

Figure 5 Data flow of the eXFlash DIMM

Architecture and components

The eXFlash DIMM has the following onboard components:

- A 19 nm multi-level cell (MLC) NAND flash memory modules
- Flash controllers that implement advanced flash management and protection techniques
- A memory controller chipset that provides an interface between a physical DDR3 memory bus and solid-state storage
- A power system that protects memory write buffers from unexpected power outage

Figure 6 shows the hardware components of the eXFlash DIMM.

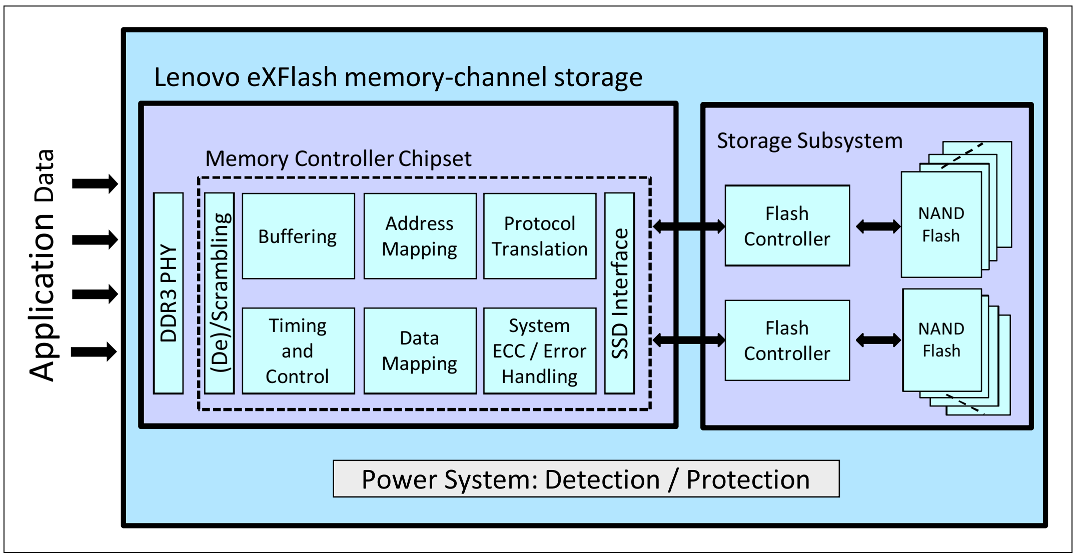

Figure 6 eXFlash DIMM hardware components

eXFlash DIMMs support the following advanced flash management technologies:

- FlashGuard, which includes innovative technologies to reliably extract more usable life from the traditional consumer-grade MLC flash than what is provided by the standard specifications that are published by NAND manufacturers.
- DataGuard, which provides full data path protection, which ensures that user data is safe throughout the entire data path. It also can recover data from failed page and NAND blocks.
- EverGuard, which prevents the loss of user data during unexpected power interruptions.

For a more information about FlashGuard technology, see *Benefits of eXFlash Memory-Channel Storage in Enterprise Solutions*, REDP-5089, which is available at this website:

http://lenovopress.com/redp5089

Supported servers

For more information about the servers that currently support eXFlash DIMMs, see the Lenovo Press Product Guide, *eXFlash DDR3 Storage DIMMs*, which is available at this website:

http://lenovopress.com/tips1141

The following eXFlash DIMM configuration rules apply:

- A minimum of one RDIMM must be installed in the same memory channel as the eXFlash DIMM.
- A maximum of one eXFlash DIMM per DDR3 memory channel is supported.
- eXFlash DIMMs support RDIMMs only; other memory types are not supported.
- eXFlash DIMMs of different capacities cannot be intermixed in the same server (for example, 200 GB and 400 GB eXFlash DIMMS cannot be intermixed).
- eXFlash DIMMs are supported only in independent memory mode. Other memory modes, such as lockstep or memory sparing, are not supported.

Table 2 lists the technical specifications for the eXFlash DIMM.

Table 2 Technical specifications for eXFlash DIMM DDR3 memory channel storage

Feature	200 GB option	400 GB option
Part number	00FE000	00FE005
Interface	DDR3 up to 1600 MHz	DDR3 up to 1600 MHz
Hot-swap device	No	No
Form factor	LP DIMM	LP DIMM
Endurance	Up to 10 drive writes per day (5-year lifetime)	Up to 10 drive writes per day (5-year lifetime)
Data Reliability	< 1 in 10^{17} bits read	< 1 in 10^{17} bits read
Shock	200 g, 10 ms	200 g, 10 ms
Vibration	2.17g rms 7-800Hz	2.17g rms 7-800Hz
Maximum power	12W	12W

For more information, see *eXFlash DIMM Configuration and Support Requirements - System x*, which is available at this website:

http://www.ibm.com/support/entry/portal/docdisplay?lndocid=SERV-FLASHDM

System configurations for eXFlash DIMM analysis

This section describes several configuration topics, including performance considerations and the evaluation tool.

Performance considerations

The performance evaluations in this paper are limited to the 400 GB eXFlash DIMMs. eXFlash DIMM performance is directly correlated with memory channel speed, processor frequency, and application load generation. Therefore, in certain situations, absolute performance values can vary by a small percentage because of a lower core count or higher processor frequency. In addition to system constraints, hardware and unique driver optimizations can affect the overall capabilities of eXFlash DIMMs.

A simple hardware optimization is to balance the memory and eXFlash configuration across the available CPU nodes. Because of memory interleaving, it is important to install a balanced configuration of RDIMMs with the eXFlash DIMMs. When the server memory configuration is not balanced, customers should expect a 25% degradation in memory performance at minimum, which also negatively affects eXFlash performance.

When eXFlash DIMMs are installed, follow the memory DIMM population rules for your Lenovo server after the RDIMMs or LRDIMMs are installed. The memory population rules ensure that the eXFlash DIMMs are evenly distributed across all CPUs for optimal performance.

Another way to optimize performance is by the use of the kernel driver. The kernel driver manages data transmission between the processor and the eXFlash DIMMs. By default, the eXFlash DIMMs are linked through their adjacent CPU nodes where each eXFlash device driver thread is assigned a 2-to-1 relationship between the thread and the respective eXFlash DIMM module. Each eXFlash DIMM device driver thread manages two eXFlash DIMM modules.

To achieve maximum performance, this relationship must be modified to a 1-to-1 relationship by using the max-occupancy option in the device driver. This relationship ensures maximum performance by allowing only the device driver to connect a single eXFlash DIMM module per processor thread. This setup is most useful when hyper-threading is enabled and the system has ample processing resources to share between applications and the device driver.

Evaluation tool

The fio tool (Flexible I/O) was used to provide the data that is presented in this paper. The tool is an accepted industry evaluation tool with low application overhead and a high level of test control. This tool was used to generate load for various workloads and to measure performance metrics, such as IOPS, latency, and bandwidth.

For this paper, fio V2.1.10 was used for all measurements.

Server configuration

Red Hat Enterprise Linux 6.5 was used for evaluation purposes. The operating system was a fresh installation, with no tuning optimizations made to the kernel to help I/O traffic or bandwidth performance. To measure low-level hardware capabilities, each eXFlash DIMM was evaluated as a raw partition in a *just a bunch of disks* (JBOD) configuration across all CPU nodes.

Early during testing, it was evident that an application must generate and request a large number of I/Os for driving each eXFlash DIMM to its full potential. Therefore, for scaling analysis, separate instances of the evaluation tool were run simultaneously to measure overall performance figures.

To ensure consistent results, all UEFI parameters that affect energy efficiency or clock frequency were disabled for generating maximum hardware performance.

All evaluations were performed by using the System x3850 X6 with four compute books, one processor in each compute book. The Intel Xeon E7-4890 v2 processors were used. In the System x3850 X6, each book contains a single multi-core processor and represents a CPU node.

As listed in Table 3, a 1 DIMM per channel (1 DPC) memory configuration consisting of eight 16 GB RDIMM at 1333 MHz was used for each book. Eight eXFlash DIMMs were populated in each CPU book across each node in a balanced configuration per the implementation guide. Table 3 also shows a snapshot of the server configuration that was used during these evaluations.

Table 3 System configuration

Configuration	Value
System configuration	
Machine Type	System x3850 X6
Processor	E7-4890 v2, 15 cores at 2.8 GHz
RAM	8x 16 GB at 1333 MHz per Book
eXFlash DIMM	32x 400 GB[a]
eXFlash DIMM slots	Each Book: slots 2, 5, 8, 11, 14, 23, 17, 23
OS	RHEL 6.5, kernel_*431
UDIMM FW	v1.5.0
FIO Version	2.1.10
UEFI Performance States Configuration	
Hyper-threading	On
Pstates	Off
C1 States	Off
C1E States	Off
Turbo	Off

a. The DIMMs that were used in this evaluation were the latest hardware revision of the product at the time of this writing. Performance characteristic can vary depending on the hardware revision that is used.

For more information about population rules, see *eXFlash DIMM Configuration and Support Requirements - System x3850/x3950 X6*, which is available at this website:

http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5096838

Methodology

The following fio parameters were changed to simulate IO patterns:

- Block size
- Queue depth
- Workload type

The following metrics were used for eXFlash DIMM performance characterization:

- IOPS, latency (µs)
- Bandwidth (MB/s)

For bandwidth analysis, we used sequential read and sequential write workloads. For IOPS and latency analysis, random read, random write, and online transaction processing (OLTP) workloads were used.

The OLTP workload I/O pattern was 67% random reads and 33% random writes. This static read/write ratio was used to gauge the performance of a mixed real-world I/O workload.

For random I/O workloads, the block size ranges were 4 K, 8 K, and 16 K.

For sequential workloads, the only block sizes used were 64 K and 1024 K. The queue depths were 1 - 128, increasing by a power of 2.

Device condition

Before each analysis, all eXFlash DIMMs were securely erased to ensure that no previous data affected performance outcomes. Then, each module was conditioned by using predefined workloads to ensure consistent steady-state performance, which is known as *preconditioning*. When eXFlash DIMMs are not preconditioned properly, results from performance evaluations can be inflated for some IO patterns and inconsistent for other IO patterns.

eXFlash DIMM performance

In this section, the performance of eXFlash is described by using different workloads, block sizes, and queue depths. All measurements are focused on the maximum hardware capabilities of a single eXFlash DIMM.

Figure 7 shows I/O performance for a single eXFlash DIMM for a 100% random read workload. As workload queue depth increases, the eXFlash DIMM achieves higher IOPs because of more available work in the transaction pipeline. For random read-intensive workloads, there is a small performance delta at low queue depths because of the lack of work queued in the pipeline.

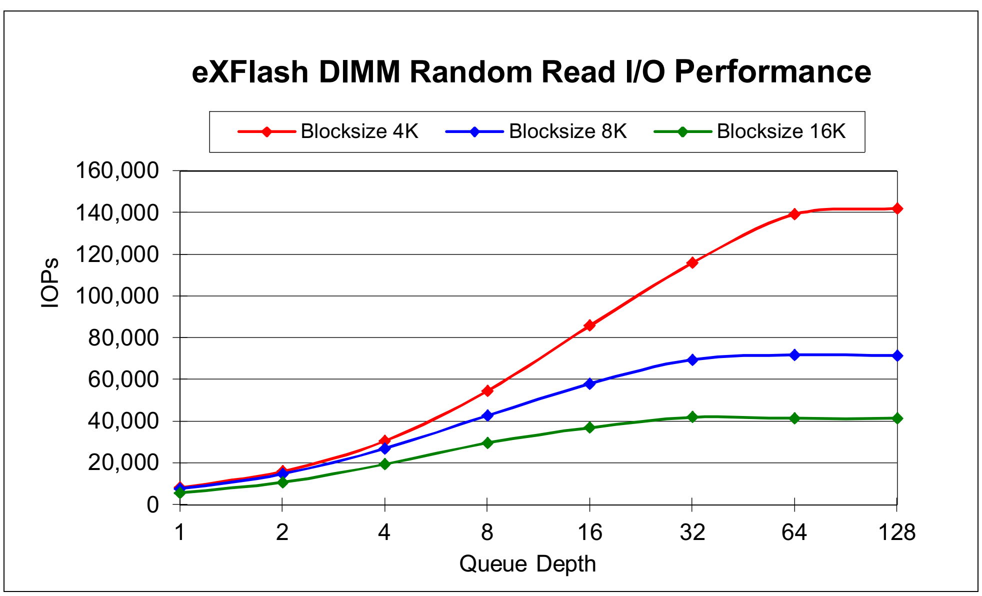

Figure 7 I/O performance for 100% random read workload

Figure 8 shows eXFlash DIMM latency performance for a 100% continuous random read-intensive workload. These latency numbers represent the time interval between fio tool I/O transaction submission and completion. As workload queue depth increases, the respective eXFlash DIMM latency increases because of the work that is waiting in the pipeline. The inflection point occurs at approximately a queue depth of 16. Therefore, the best trade-off between maximum I/O performance and the best latency numbers is approximately this queue depth.

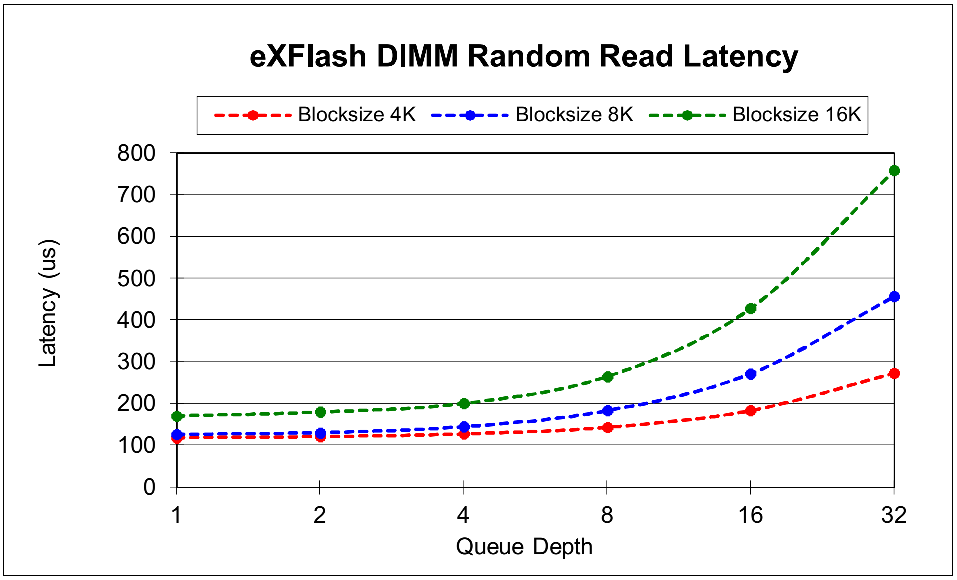

Figure 8   Latency performance for 100% random read workload

Figure 9 shows the eXFlash DIMM I/O performance for a 100% continuous random write-intensive workload. The random write IOPs of an eXFlash DIMM is negligibly affected by increasing the queue depth. The IOPs reduction with larger block sizes is because of reaching the peak bandwidth capability of the eXFlash DIMM. It is expected that fewer IOPs are processed as the block size increases.

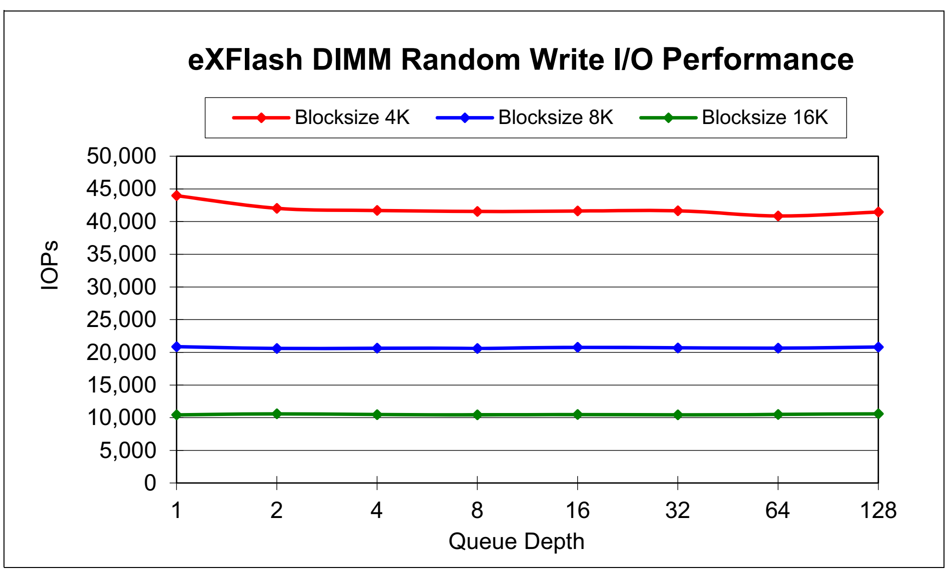

Figure 9   I/O performance for 100% continuous random write workload

Figure 10 shows eXFlash DIMM latency performance for a 100% continuous write-intensive workload without any pauses between transactions. These results demonstrate the maximum hardware capabilities of the eXFlash DIMM when application overhead is removed. Best case latency numbers are listed in Table 5 on page 19.

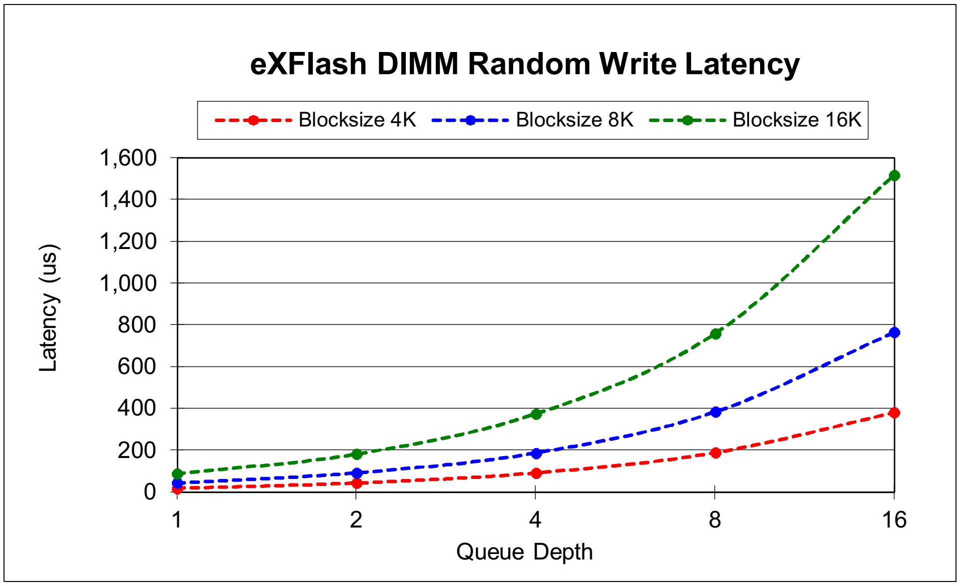

Figure 10 Latency performance for 100% continuous random write workload

The WriteNow feature of the eXFlash DIMMs reduces write latency by allowing data to be written to flash memory at an accelerated rate. For applications that use small block sizes at low queue depths, performance improvements can be up to 70% with write latency measuring as low as 5 µs. This feature is most effective when the transaction stream varies between reads and writes, or when brief application pauses occur in the write data stream.

To examine this latency performance, the fio ThinkTime parameter was used to insert a delay between write operations for a static workload and queue depth.

Figure 11 shows eXFlash DIMM latency performance for a 100% random write-intensive workload with a varying amount of fio ThinkTime between write transactions. As the ThinkTime parameter is increased to 16 µs, there is a small (7%) reduction in I/O, but a significant (71%) reduction in write latency. A single eXFlash DIMM demonstrated 5.35 µs of latency for a random write workload that uses a 4 K block size at a queue depth of 1 when the write transactions were separated by a 16 µs interval in the data stream.

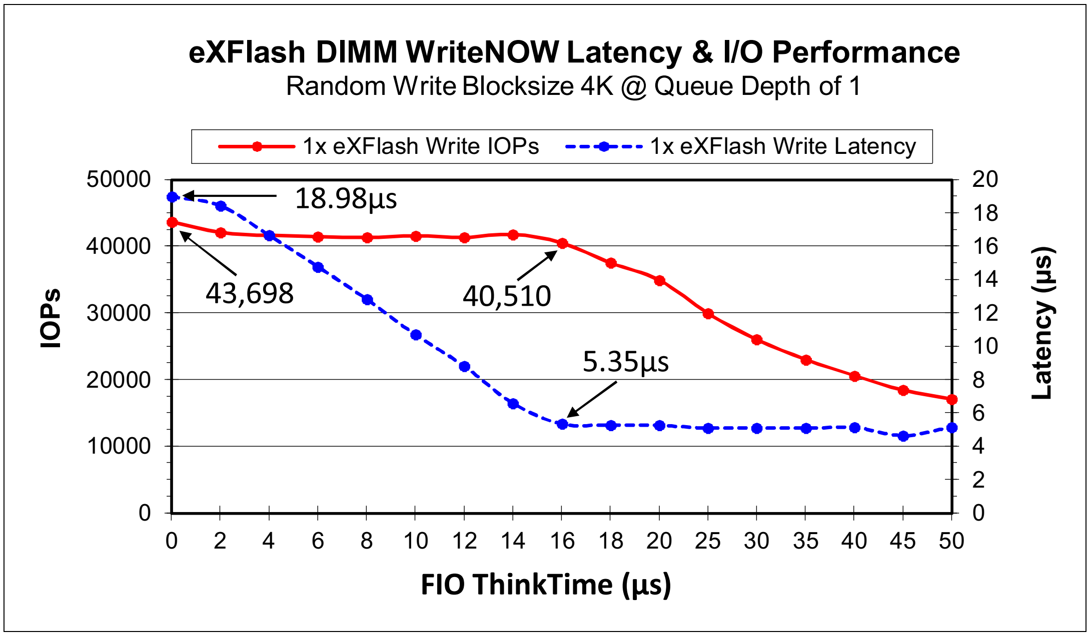

Figure 11 WriteNow I/O and latency performance for 100% random write workload

Figure 12 shows eXFlash DIMM I/O performance for an OLTP workload I/O pattern of 67% random reads and 33% random writes. For an OLTP workload, there is a greater I/O performance separation between the transaction block sizes because of the nature of the mixed workload.

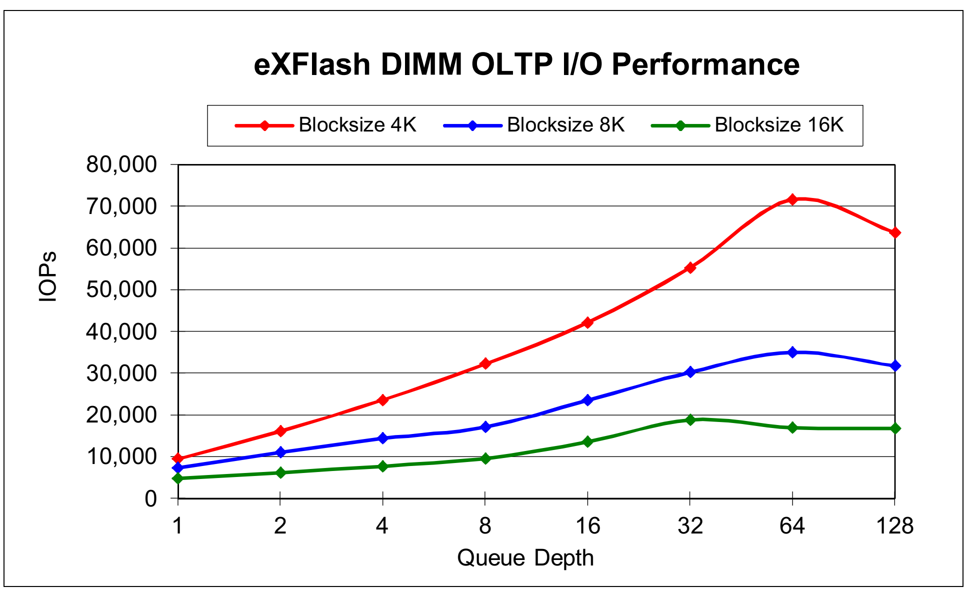

Figure 12 I/O performance for OLTP workload using 67/33 random read/write ratio

Figure 13 shows eXFlash DIMM latency performance for an OLTP workload I/O pattern of 67% random reads and 33% random writes. These results demonstrate the maximum hardware capabilities of the eXFlash DIMM when application overhead was removed. Our best case latency numbers are listed in Table 4 on page 19.

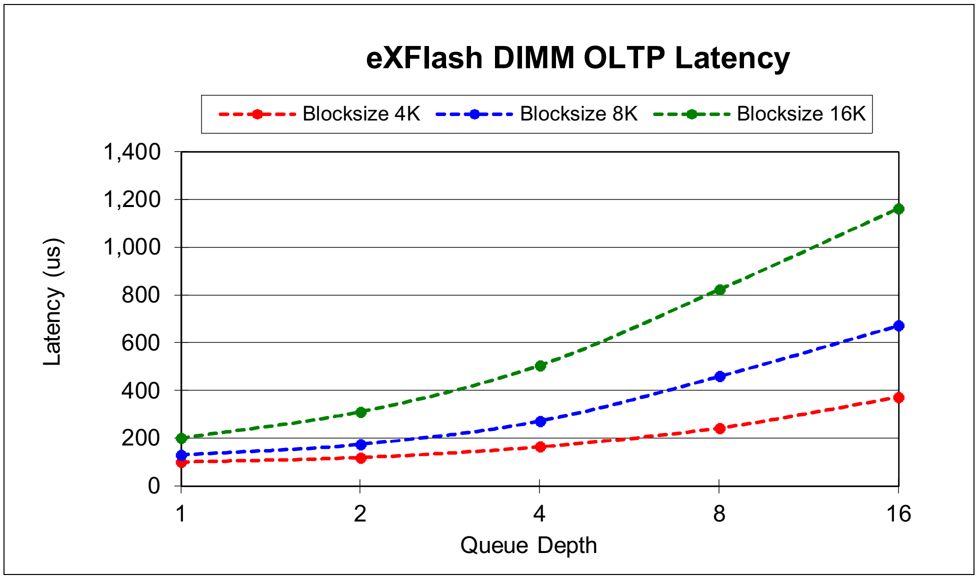

Figure 13 Latency performance for OLTP workload using 67/33 random read/write ratio

Figure 14 shows eXFlash DIMM bandwidth performance for a 100% sequential read and 100% sequential write workload. For sequential read workloads, the larger block sizes generate maximum bandwidth performance at smaller queue depths because of the amount of information that is transferred. For sequential write workloads, bandwidth performance is negligibly affected by transaction block size and does not vary by queue depth, so the same bandwidth is measured with 64 K and 1024 K block sizes.

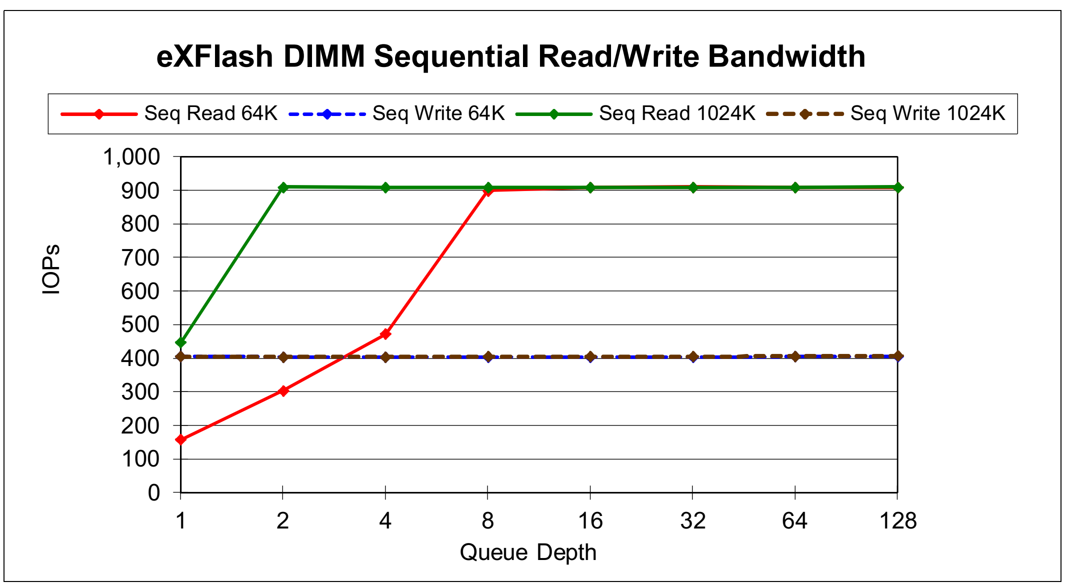

Figure 14 Bandwidth performance for sequential read and write workloads

eXFlash scaling performance

Figure 15 shows eXFlash DIMM I/O performance for multiple DIMMs scaled across CPU books in a System x3850 X6. The values in Figure 15 were normalized based on the results of a single eXFlash DIMM. A 4 K block size was used, but similar scaling results are expected for other block sizes. For scaling evaluations, individual instances of fio operated on each eXFlash DIMM to drive enough I/O to measure maximum performance.

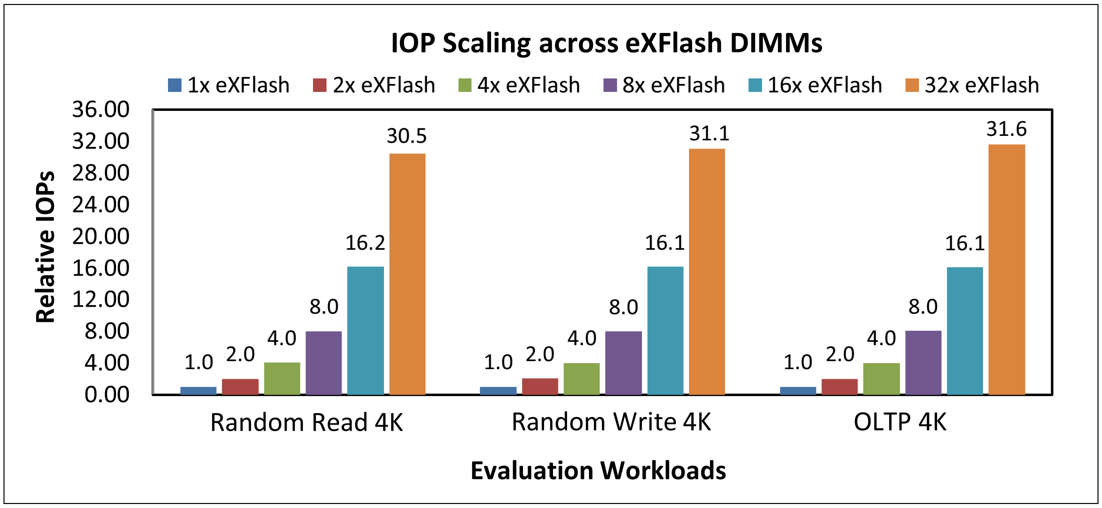

Figure 15 Normalized I/O scaling performance for eXFlash DIMMs

Figure 16 shows eXFlash latency performance when multiple eXFlash DIMMs were exercised across CPU books in the system. The measured values in Figure 16 are the average result across each grouping of scaled eXFlash DIMMs. A 4 K block size produced the best result; therefore, this block size was used for comparison. For the random write workload, no delays are inserted between write transactions. eXFlash DIMM performance scaled linearly up to 16 DIMMs. The performance at 32 DIMMs was slightly less than linear because of the CPU cores becoming fully used by the fio tool and the eXFlash DIMM device driver threads.

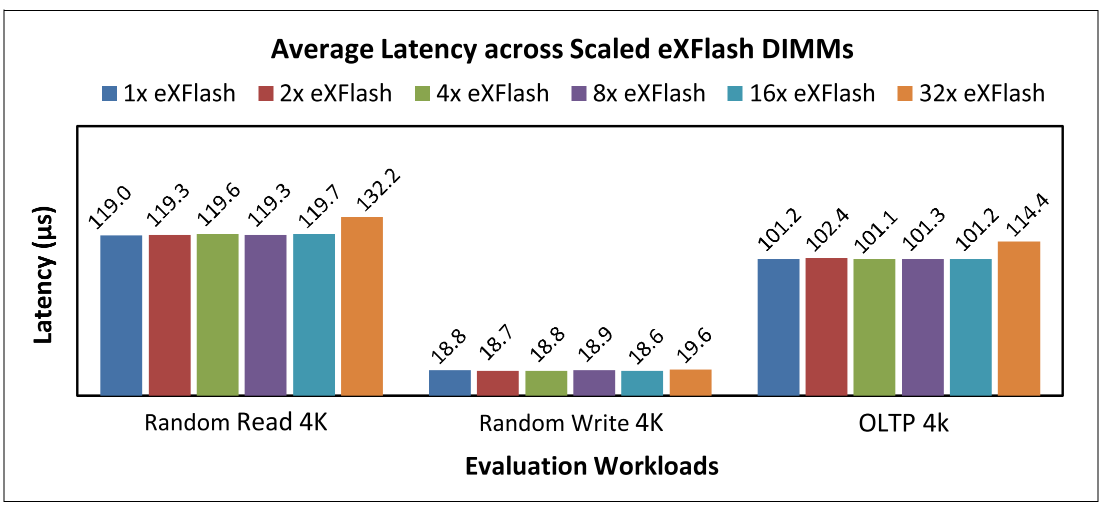

Figure 16 Average latency performance across scaled eXFlash DIMMs

Figure 17 shows eXFlash bandwidth performance when multiple eXFlash DIMMs are exercised across CPU nodes in the system. The values in Figure 17 are normalized results, based on the value that is measured on a single eXFlash DIMM. For these workloads, the measured linear scaling performance remains independent of block size. Absolute values are listed in Table 4 on page 19.

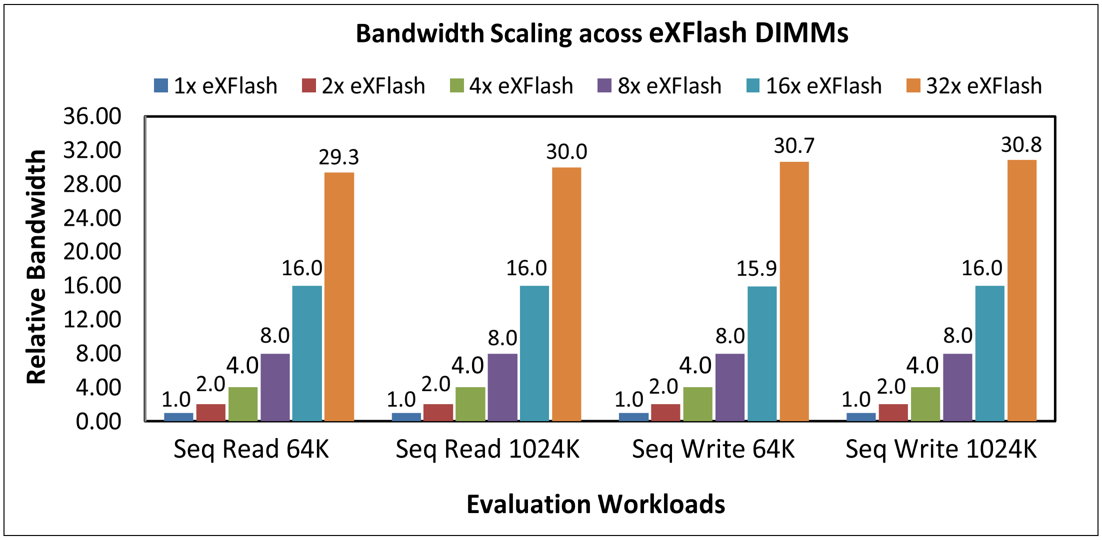

Figure 17 Normalized bandwidth performance for eXFlash DIMMs

Conclusion

The measurements that are presented in this paper provide a snapshot of the capabilities of Lenovo's eXFlash DIMM memory-channel storage products. eXFlash DIMMs can add capacity to a server with limited space or provide ultra-fast caching.

The linear nature of eXFlash DIMM scaling also enhances I/O and bandwidth performance across DIMMs without experiencing a significant drop in performance. Lenovo's eXFlash DIMM memory channel storage provides you with a compact, low-maintenance, high-performance device.

Lenovo memory-channel is another option to consider when looking at increasing storage performance in supported System x and Flex System servers.

Appendix

The tables that are in this appendix list the data that was gathered during our performance evaluation. Table 4 lists the IOPS data points and Table 5 lists the bandwidth and latency performance figures.

Table 4 IOPS performance for 400 GB eXFlash DIMMs: Maximum values recorded during scaling analysis

Block size	IOPS performance					
	1x eXFlash DIMM	**2x eXFlash DIMMs**	**4x eXFlash DIMMs**	**8x eXFlash DIMMs**	**16x eXFlash DIMMs**	**32x eXFlash DIMMs**
Random Read (IOPS)						
4k block	142K	286K	574K	1,141K	2,299K	4,331K
8k block	72K	145K	291K	577K	1,163K	2,163K
16k block	42K	84K	168K	336K	673K	1,286K
Random Writes (IOPS)						
4k block	44K	89K	176K	351K	710K	1,366K
8k block	21K	42K	83K	166K	334K	670K
16k block	11K	21K	43K	85K	171K	342K
OLTP Read/Write (IOPS)						
4k block	72K	144K	288K	576K	1,154K	2,264K
8k block	35K	71K	142K	284K	568K	1,136K
16k block	19K	37K	75K	150K	303K	586K

Table 5 Bandwidth performance for 400 GB eXFlash DIMMs: Maximum values recorded during scaling analysis

Workload	Bandwidth and latency performance					
	1x eXFlash DIMM	**2x eXFlash DIMMs**	**4x eXFlash DIMMs**	**8x eXFlash DIMMs**	**16x eXFlash DIMMs**	**32x eXFlash DIMMs**
Sequential Throughput Bandwidth (MB/s) @ 64K block size with Queue Depth = 16						
100% Seq Read	909 MBps	1,819 MBps	3,639 MBps	7,278 MBps	14,556 MBps	27,908 MBps
100% Seq Write	405 MBps	810 MBps	1,620 MBps	3,240 MBps	6,479 MBps	12,130 MBps
Sequential Throughput Bandwidth (MB/s) @ 1024K block size with Queue Depth = 16						
100% Seq Read	909 MBps	1,815 MBps	3,632 MBps	7,252 MBps	14,483 MBps	27,869 MBps
100% Seq Write	406 MBps	812 MBps	1,623 MBps	3,248 MBps	6,499 MBps	12,518 MBps
Latency (μs) @ 4K block size with Queue Depth = 1						
100% Random Read	119.01 μs	119.34 μs	119.58 μs	119.34 μs	119.65 μs	132.24 μs
100% Random Write	18.84 μs	18.66 μs	18.93 μs	18.74 μs	18.61 μs	18.70 μs
OLTP (67% R / 33% W)	101.18 μs	102.44 μs	101.13 μs	101.35 μs	101.18 μs	114.38 μs

Related publications

The following publications provide more information about the topic in this document. Some of the publications that are referenced in this list might be available in softcopy only:

- *Benefits of eXFlash Memory-Channel Storage in Enterprise Solutions*, REDP-5089

  http://lenovopress.com/redp5089
- *eXFlash DDR3 Storage DIMMs*, TIPS 1141

  http://lenovopress.com/tips1141
- eXFlash DIMM Configurations and Support Requirements - System x.

  http://www.ibm.com/support/entry/portal/docdisplay?lndocid=SERV-FLASHDM
- *Lenovo X6 Servers: Technical Overview*, REDP-5059

  http://lenovopress.com/redp5059
- *Workload Optimization with IBM X6 Servers*, REDP-5058

  http://lenovopress.com/redp5058
- *System x3850 X6 Planning and Implementation Guide*, SG24-8208

  http://lenovopress.com/sg248208

About the author

Tristian “Truth” Brown is a Hardware Performance Engineer on the Lenovo EBG Server Performance Team in Raleigh, NC. He is responsible for the hardware analysis of high-performance, flash-based storage solutions for System x servers. Truth earned a bachelor’s degree in Electrical Engineer from Tennessee State University and a master’s degree in Electrical Engineering from North Carolina State University. His focus areas were in Computer Architecture and System-on-Chip (SoC) microprocessor design and validation.

Thanks to the following people for their contributions to this project:

David Watts
Senior IT Consultant
Lenovo Press

Karen Lawrence
Technical Writer
IBM Redbooks

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document REDP-5188-00 was created or updated on May 21, 2015.

Send us your comments in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:
  ibm.com/redbooks
- Send your comments in an email to:
  redbooks@us.ibm.com

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

eXFlash™
Flex System™
Lenovo®
Lenovo(logo)®
System x®

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.