Performance Brief

Using Data Deduplication and Compression with KIOXIA CM6 Series SSDs to Accelerate VMware® Storage in Storage Arrays

Introduction

One of the most common deployment methods for data center storage is using storage arrays. A storage array provides a way to quickly access and store mission-critical data, virtual machines (VMs), large databases and other workloads. They deliver high performance, high storage density and the ability to provision storage specifically to the requirements of the workload itself. They feature advanced capabilities that may include redundant array of independent disks (RAID), virtualization, encryption, data deduplication and compression.

Data deduplication (dedupe) and compression in particular can significantly improve SSD value as they enable a host to store more data on the available storage capacity. Dedupe is a form of compression that eliminates duplicate copies and other repeating data. Instead of using a compression algorithm, the storage array looks for duplicate data patterns. When a data pattern is found, only one instance of that data is written to the drive and files are located using pointers. Dedupe enables the usable storage capacity to be increased to accommodate specific application workloads.

Compression compacts data before a page of data is written to flash memory. When storing or transmitting large data sizes, compression is a common method. In a storage array, the server writes data onto a volume and runs a compression algorithm to reduce the amount of physical space required on the array side.

When both dedupe and compression are enabled, they create capacity savings and can improve storage utilization that can lower capital expenditures as the number of drives required to meet capacity needs can be reduced. They can also lower the overall system cost per gigabyte[1] (GB) and the cost per Input/Output (I/O) operation. However, these capabilities require additional CPU resources to remove duplicate data blocks and to compress the remaining data that can lead to performance degradation.

This performance brief presents a comparison of virtual machines running Microsoft® SQL Server™ database storage performance with dedupe and compression (enabled and disabled) on a storage array deployed with twenty-eight KIOXIA CM6-R Series enterprise NVMe™ SSDs. This comparison determines if the storage array deployed with KIOXIA CM6-R Series SSDs can perform dedupe and compression capabilities without experiencing a performance hit.

The tests utilized HammerDB[2] software and TPROC-C[3] online transaction processing (OLTP) workloads to measure database throughput (transactions per minute), virtual machine CPU utilization and DRAM usage. Prior to the tests, the optimal compression rate for the test database was measured at 6.7:1.

The test results show that with dedupe and compression enabled, the test system did not experience significant performance degradation – database throughput, CPU utilization and DRAM usage test results were comparable to the disabled dedupe and compression system.

The results presented include a brief description of each workload test, a graphical depiction of the test results and an analysis. Appendix A covers the hardware and software configuration – Appendix B covers the configuration set-up and test procedures.

Test Results Snapshot

For VMware ESXi™ 8.0, a 6.7:1 data compression rate of virtual machines delivered ~96% TPM efficiency

Database Throughput

With Dedupe-Compression Disabled: ***1,309,401 TPM***

With Dedupe-Compression Enabled: ***1,258,899 TPM***

~4% Difference
~96% TPM efficiency

Virtual Machine CPU Utilization

With Dedupe-Compression Disabled: ***16.4%***

With Dedupe-Compression Enabled: ***16.8%***

~0.4% Difference
Negligible CPU resources used

DRAM Usage

With Dedupe-Compression Disabled: ***58.1 GB***

With Dedupe-Compression Enabled: ***58.2 GB***

~0.1 GB Difference
Negligible system memory used

Test Results[4]

Compression Rate

Before testing, the storage array measured various compression rates to determine the optimal rate to use for the 580 GB database created by HammerDB test software. The compression rate is a ratio of the 'used before space' divided by the 'used after space.' If the 'used before space' is 3 GB and the 'used after space' is 1 GB, for example, then the compression rate would be 3:1 (or 3x).

During testing, the compression rate reached an 11:1 ratio for the 580 GB database size, however, the stabilized compression rate was 6.7:1. Depending on the data type, the compression rate will vary. Some data, such as video and other very large files are functionally uncompressible. Virtual machines (especially clones) are compressible, as are many databases. For this comparison, virtual machine clones were created, and initially, they delivered high compression. When additional random data was run on the clones, compression was more difficult to achieve, thus lowering the rate. The compression rate remained steady at 6.7:1 even when excessive random data was written to the virtual machine clones.

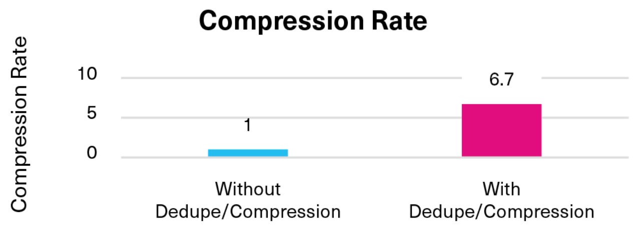

Database Throughput / TPM

This test measured the number of SQL Server transactions completed per minute on average for each virtual machine on both configurations. HammerDB software executed the TPROC-C transaction profile that randomly performed transaction types such as new orders, payment, order status, delivery and stock levels, simulating an OLTP environment. This test case represented a large number of users that conduct simple and short transactions where sub-second response times are required. The test results were recorded in transactions per minute (TPM) - the higher result is better.

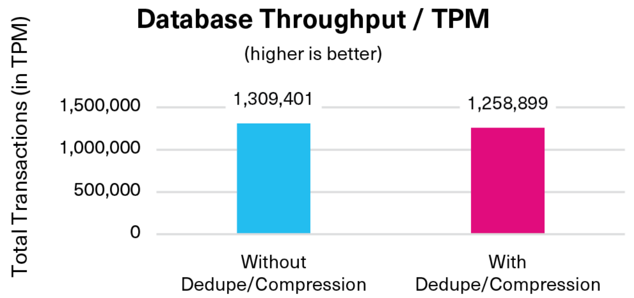

Virtual Machine CPU Utilization

This test measured the percentage of the virtual machine CPU cycles that were being used for a given database workload to ensure that the CPUs were efficiently processing them. The base number for this CPU utilization was measured at 16.4% for the workload performed by the database server against the storage array when dedupe and compression were disabled. With dedupe and compression enabled, CPU utilization was 16.8%.

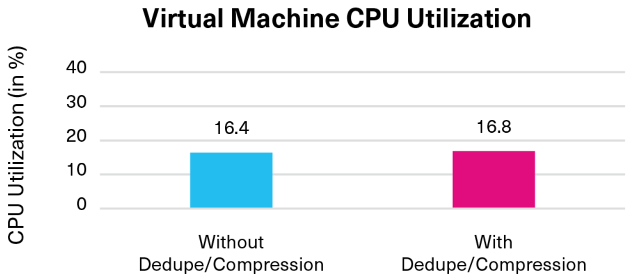

DRAM Usage

This test measured the amount of DRAM (system memory) used by the virtual machines. The base number for DRAM usage was measured at 58.1 GB for the workload performed by the database server against the storage array when dedupe and compression were disabled. With dedupe and compression enabled, DRAM usage was measured at 58.2 GB.

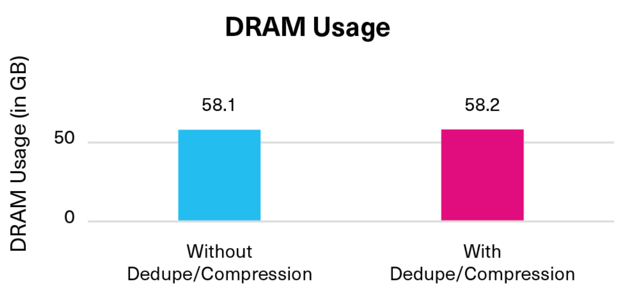

Test Analysis

Without dedupe and compression enabled, the test system delivered 1,309,401 TPM. With dedupe and compression enabled, the system delivered 1,258,899 TPM. The margin of deviation between these results was about 4%, representing about 96% performance efficiency with dedupe and compression enabled. The high performance delivered by KIOXIA CM6 Series SSDs helped to accelerate the SQL Server database in both test cases.

Regarding virtual machine CPU utilization and DRAM usage, the test results show that even with dedupe and compression enabled, the test system only required about 0.4% more CPU resources and 0.1 GB of additional system memory. Based on these test results, dedupe and compression should always be enabled, and in doing so, more storage capacity from the SSDs in the storage array can be obtained.

By dynamically provisioning volumes, creating shared volumes and clusters, and using deduplication and compression, less drive space is required on storage arrays, improving operational efficiency and total cost of ownership.

Summary

Storage arrays are a common way to deploy virtual machines for most enterprises. When performing dedupe and compression, some storage arrays experience performance degradation. With KIOXIA CM6 Series SSDs deployed in the storage array, the test system demonstrated negligible performance degradation when enabling these capabilities without using valuable CPU resources or additional system memory. The ability to centralize data enables high deduplication rates.

Conversely, for other types of storage such as scale-out or VMware vSAN™, the overall compression rate is generally not as high as storage arrays because of the replication factors associated with data redundancy, as well as the difficulty of deduplicating data stored on different nodes. Many scale-out or hyperconverged solutions also require additional resource utilization so that each host can handle replication and data resiliency, which in turn can require additional nodes to handle the workload requirements. However, a dedicated storage array means that host side resources can be dedicated for VM usage. Additionally, by compressing the storage array data into a smaller footprint, the overall solution can provide cost of ownership savings and maximum use of storage capacity while delivering good end user experiences. The KIOXIA CM6 Series SSDs and the storage array delivered dedupe and compression capabilities without a significant performance hit.

KIOXIA CM6-R Series Enterprise NVMe SSD Information

KIOXIA CM6-R Series enterprise NVMe SSDs for read intensive applications support the PCIe® 4.0 interface, 1 DWPD[5] endurance, capacities up to 30.72 terabytes[1] (TB), dual-port capabilities for high availability, up to 21-watt power envelope, and a host of security options[6] – all geared to support a wide variety of workload requirements.

More information.

KIOXIA CM6 Series SSD[7]

Appendix A

Hardware/Software Test Configuration

Server Information	
Number of Servers	3
Number of CPU Sockets	2
CPU	AMD EPYC™ 7702
Number of CPU Cores	64
CPU Frequency	2.0 gigahertz (GHz)
Total Memory	512 GB DDR4 DRAM
Memory Frequency	DDR4-3200
Server HBA	Marvell® FastLinQ™ QL41000
HBA Speed	32 gigabits per second (Gb/s)

SSD Information	
Model	KIOXIA CM6-R Series
Interface	PCIe 4.0 x4
Number of Devices	28
Form Factor	2.5-inch[8] (U.3)
Capacity	3.84 TB
DWPD	1 (5 years)
Active Power	up to 19 watts

Operating System Information	
Operating System	VMware ESXi™
Version	8.0.1

Storage Array Information	
Operating System Version	9.5.11
Number of Controllers	4
Number of Drive Enclosures	2
Number of SSDs	28
Number of Ports	48
Total Memory	768 gibibytes[9] (GiB)
HBA Port Speed	32 Gb/s
Fibre Channel Ports Used	8 (2 per Controller)

Database Software Information	
Database Software Model	Microsoft® SQL Server™ Enterprise (64-bit)
Version	16.0.1050.5
Operating System	Microsoft Windows® Server® 2022

VM Test Information	
Number of VMs	12
Number of VM CPUs	16
VM DRAM Size	16 GB
VM Disk Size	1 x 100 GB OS vDisk 1 x 1,000 GB vDisk SQL

Test Software Information	
Test Software	HammerDB TPROC-C
Version	4.5
Number of Virtual Users	256

Appendix B

Configuration Set-up/Test Procedures

Configuration Set-up

The test system was configured using the hardware and software equipment outlined in Appendix A.

The test server was connected to a Fibre Channel (FC) switch via the 32 Gb/s Marvell FastLinQ FC host bus adapter.

The storage array was connected to the same FC switch as the test server.

The test server was installed with Windows Server 2022.

Twenty-eight KIOXIA 3.84 TB CM6-R Series SSDs were deployed into the front bay of the storage array.

VMware ESXi was installed on all four controllers of the storage array.

The storage array was configured with two volumes:

- One 16 TB volume with dedupe and compression turned on (Enabled)
- One 16 TB volume with dedupe and compression turned off (Disabled).

Both volumes were presented over the FC fabric to the three VMware cluster nodes.

The Volumes were formatted with VMFS6[10] to be used as virtual machine storage.

Twelve Microsoft Windows virtual machines were created, each with Microsoft Windows Server 2022 and Microsoft SQL Server 2022 -- each virtual machine included 100 GB operating system disk space and 1,000 GB SQL data disk space.

Sixteen virtual CPUs and 64 GB of DRAM was added to each of the twelve Microsoft Windows virtual machines.

HammerDB was installed to create a 580 GB test database on the disabled 16 TB dedupe and compression volume.

The test database was backed up and restored on both dedupe and compression volumes to ensure the databases were identical when the tests were run.

HammerDB software was tested with different user counts to determine the optimal number of virtual users. The test iterations determined that two hundred and fifty-six virtual users provided the optimal user count.

Test Procedures

The tests were run utilizing HammerDB software and two hundred and fifty-six virtual users against the identical 580 GB databases stored on both dedupe and compression volumes (one enabled/one disabled).

The tests were run on the 3.84 TB KIOXIA CM6-R Series SSDs and the test results were recorded for database throughput (in TPM), virtual machine CPU utilization (in percentages) and DRAM utilization (in GB size).

NOTES:

[1] Definition of capacity - KIOXIA Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes, a terabyte (TB) as 1,000,000,000,000 bytes and a petabyte (PB) as 1,000,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1Gbit = 2^{30} bits = 1,073,741,824 bits, 1GB = 2^{30} bytes = 1,073,741,824 bytes, 1TB = 2^{40} bytes = 1,099,511,627,776 bytes and 1PB = 2^{40} bytes = 1,125,899,906,842,624 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

[2] HammerDB is open source database benchmarking and load testing software. It identifies transactional use cases and provides meaningful information about the data environment that can include performance comparisons.

[3] TPROC-C is the OLTP test workload implemented in HammerDB benchmarking software derived from the TPC-C™ Benchmark Standard.

[4] Read and write speed may vary depending on the host device, read and write conditions and file size.

[5] Drive Write(s) per Day: One full drive write per day means the drive can be written and re-written to full capacity once a day, every day, for the specified lifetime. Actual results may vary due to system configuration, usage, and other factors.

[6] Optional security feature compliant drives are not available in all countries due to export and local regulations.

[7] The product image shown is a representation of the design model and not an accurate product depiction.

[8] 2.5-inch indicates the form factor of the SSD and not the drive's physical size.

[9] One gibibyte (GiB) equals 2^{30} or 1,073,741,824 bytes.

[10] VMware vSphere® VMFS functions as a volume manager, enabling VM files to be stored in logical containers called VMFS datastores. VMFS6 represents the version used in the comparison.