# Fujitsu Server PRIMERGY & PRIMEQUEST Performance Report RAID Controller Performance 2021

This document is intended for those responsible for disk IO performance on Fujitsu PRIMERGY and PRIMEQUEST servers. It provides performance information to help you understand the various RAID controller options and application areas that apply to your internal disk subsystem. The recommended controller selection and parameter settings depend on your data safety and performance requirements, as well as your server configuration. This document introduces the current generation of controllers that can be installed in PRIMERGY and PRIMEQUEST systems since 2021.

| Version |
| --- |
| 1.5<br>2024-07-16 |


 

# Contents



 

 

 

# Introduction

In a server environment, internal storage is an important component that affects system performance. It is also necessary to protect data from hardware failures and to have a mechanism that can be easily recovered. Therefore, the server adopts a method called RAID (Redundant Array of Independent Disks), which improves performance by controlling and accessing multiple drives, and provides the ability to complement the data by the remaining drives in the event of a drive failure. PRIMERGY and PRIMEQUEST servers are available in various configurations with various RAID controllers and drives.

The purpose of this document is to summarizes the performance data of major storage drives connected with RAID controllers that are newly supported by PRIMERGY and PRIMEQUEST with the 2nd generation or later Intel Xeon scalable processor, and to refer to it when constructing a storage system. For the functions of the controller and each setting value, refer to the User Manual of each controller. See the White Paper below for other RAID controllers.

- "RAID Controller Performance 2018"
  Describes the RAID controllers provided for PRIMERGY and PRIMEQUEST, including their performance.
- "Disk I/O Performance Basics"
  Provides an overview of disk I/O performance.

In this document, a power of 10 (1 TB = $10^{12}$ bytes) is used to indicate the capacity of the storage drive, and a power of 2 (1 MiB/s = $2^{20}$ bytes/s) is used to indicate all other capacities, file size, block size, and throughput.

 

# Measurement Target RAID Controllers

Performance measurements were performed by connecting various storage drives to the following RAID controllers.

| Controller name | Mounting method | Cache | Supported interface | | Maximum number of drives | RAID level |
|---|---|---|---|---|---|---|
| | | | CPU side | Driver side | | |
| Intel VROC SATA RAID Controller | RX2540 M6 Onboard | - | DMI3 x4 | SATA 6G | 8 | 0/1/10 |
| Intel VROC NVMe RAID Controller | RX2540 M6 Onboard | - | PCIe Gen4 x16 | NVMe Gen4 | 8 | 0/1/10/5 |
| Embedded MegaRAID [Intel C620 PCH SATA Controller] | RX2540 M5 Onboard | - | DMI3 x4 | SATA 6G | 8 | 0/1/10 |
| Microsoft storage pool [Intel PCIe Device Adapter] | RX2540 M5 Onboard | - | PCIe Gen3 x16 x2 | NVMe Gen3 | 8 | 0/1/10 |
| PSAS CP2100-8i | PCIe Option | - | PCIe Gen3 x8 | SATA 6G SAS 12G | 8 | 0/1/10/5 |
| PSAS CP2200-16i | PCIe Option | - | PCIe Gen4 x8 | SATA 6G SAS 24G | 16 | 0/1/10/5 |
| | | | | NVMe Gen4 | 4 | |
| PRAID CP400i | PCIe Option | - | PCIe Gen3 x8 | SATA 6G SAS 12G | 8 | 0/1/1E/10/5/50 |
| PRAID CP500i | PCIe Option | - | PCIe Gen3 x8 | SATA 6G SAS 12G | 8 | 0/1/10/5/50 |
| PRAID CP600i | PCIe Option | - | PCIe Gen4 x8 | SATA 6G SAS 12G | 8 | 0/1/10 |
| PRAID EP400i | PCIe Option | 1 GB | PCIe Gen3 x8 | SATA 6G SAS 12G | 8 | 0/1/1E/10/5/50 /6/60 |
| PRAID EP420i | PCIe Option | 2 GB | PCIe Gen3 x8 | SATA 6G SAS 12G | 8 | 0/1/1E/10/5/50 /6/60 |
| PRAID EP540i | PCIe Option | 4 GB | PCIe Gen3 x8 | SATA 6G SAS 12G | 16 | 0/1/1E/10/5/50 /6/60 |
| | | | | NVMe Gen3 | 4 | |
| PRAID EP640i | PCIe Option | 4 GB | PCIe Gen4 x8 | SATA 6G SAS 12G | 8 | 0/1/1E/10/5/50 /6/60 |
| PRAID EP680i | PCIe Option | 8 GB | PCIe Gen4 x8 | SATA 6G SAS 12G | 16 | 0/1/1E/10/5/50 /6/60 |
| | | | | NVMe Gen4 | 4 | |
| PRAID EP3254-8i | PCIe Option | 4 GB | PCIe Gen4 x8 | SATA 6G SAS 24G | 8 | 0/1/10/5/50 /6/60 |
| PRAID EP3258-16i | PCIe Option | 8 GB | PCIe Gen4 x8 | SATA 6G SAS 24G | 16 | 0/1/10/5/50 /6/60 |
| | | | | NVMe Gen4 | 4 | |

 

# Controller Interfaces and Their Throughput Limits

The RAID controller has an interface that connects to the storage drive and an interface that connects to the CPU. The drive interfaces are Serial Attached SCSI (SAS), Serial Advanced Technology Attachment (SATA), and Non-Volatile Memory Express (NVMe). NVMe is a communication protocol designed for flash storages that use Peripheral Component Interconnect Express (PCIe) for connectivity. A typical NVMe drive uses 4 lanes PCIe. The interface of the CPU is PCIe, but in the case of the onboard SATA controller, it is connected by Direct Media Interface (DMI).

The throughput upper limits of the storage drive side interface and the CPU side interface are as follows.

### Storage drive side

| Interface | Communication band | Theoretical throughput | Effective throughput (90%) |
|---|---|---|---|
| SATA 6G | 6 Gbps | 572 MiB/s | 515 MiB/s |
| SAS 12G | 12 Gbps | 1,144 MiB/s | 1,030 MiB/s |
| NVMe Gen3 | 8 Gbps x4 | 3,756 MiB/s | 3,380 MiB/s |
| NVMe Gen4 | 16 Gbps x4 | 7,512 MiB/s | 6,760 MiB/s |

The theoretically achievable throughput of SATA and SAS is calculated by subtracting 20% of serial transfer redundancy with 8b/10b coding. For NVMe, it subtracts 1.54% redundancy with 128b/130b coding. The actual achievable throughput can be estimated by multiplying this value by 0.90. This 90% value is an average empirical value calculated from the values of various components that have been measured in the past.

The maximum throughput on the drive side is the throughput of the interface used for connection times the number of drives. You can use the expander to expand the number of connected drives but it doesn't increase throughput because it only distributes the data flow. For example, even if 16 drives are connected to a controller with 8 ports of SAS 12G using an expander, the upper limit is 8,240 MiB/s (1,030 MiB/s x 8 ports).

### CPU side

| Interface | Number of lanes | Communication band | Theoretical throughput | Effective throughput (90%) |
|---|---|---|---|---|
| DMI Gen3 | x4 | 8 Gbps x4 | 3,756 MiB/s | 3,380 MiB/s |
| PCIe Gen3 | x8 | 8 Gbps x8 | 7,512 MiB/s | 6,760 MiB/s |
| PCIe Gen3 | x16 | 8 Gbps x16 | 15,024 MiB/s | 13,520 MiB/s |
| PCIe Gen4 | x8 | 16 Gbps x8 | 15,024 MiB/s | 13,520 MiB/s |
| PCIe Gen4 | x16 | 16 Gbps x16 | 30,048 MiB/s | 27,040 MiB/s |

The theoretically achievable throughput is calculated by subtracting 1.54% redundancy with 128b/130b coding. The actual achievable throughput can be estimated by multiplying this value by 0.90. This 90% value is an average empirical value calculated from the values of various components that have been measured.

 

The maximum throughput that a RAID controller can actually achieve is either the maximum throughput on the drive side or the maximum throughput on the CPU side, whichever is smaller. For example, if you connect eight SAS 12G storage drives to a RAID controller, the drive side throughput is expected to be 8,240 MiB/s (8 x 1,030 MiB/s), but if the CPU side is PCIe Gen3 x8, the throughput is limited to only 6,760 MiB/s. Similarly, when 4 NVMe Gen3 storage drives are connected, the drive-side throughput is expected to be 13,520 MiB/s (4 x 3,380 MiB/s), but the CPU side is limited to 6,760 MiB/s for PCIe Gen3 x8. When adopting a RAID controller, it is necessary to consider the type of the storage drive you want to use, the number of drives, and the type of the interface on the CPU side.

## *RAID Level and the Expected Performance*

RAID level depends on the control method of the connected drive. The levels are defined between RAID0 and RAID6. The commonly used levels are as follows.

**RAID0 Striping**

You can increase the performance by distributing data to multiple drives for reading and writing. It has no redundancy and all data will be lost if one drive fails. Therefore, the failure rate will be higher by the number of connected devices.

**RAID1 Mirroring**

By writing the same data to two drives, the system can continue to operate in the event of a drive failure by using the remaining drive. Read performance can be improved by accessing two drives in parallel, but write performance will be degraded because the same data is written to two drives.

**RAID10**

By combining RAID1 (mirroring) for high reliability and RAID0 (striping) for high speed, access performance and fault tolerance can be improved. Write performance can also be improved by increasing the number of mirroring sets.

**RAID5**

It distributes data to multiple drives and at the same time generates and writes parity data. Parity data is distributed and written to all drives. By regenerating data from parity data, it can have tolerance of a single drive failure. Uses the capacity of one connected drive for parity data. Therefore, the capacity of the number of drives minus one can be used.

**RAID6**

It distributes data to multiple drives and at the same time generates double parity data and writes it. By having double parity data, it can have tolerance of two drive failures. Since the capacity of two connected drives is used for parity data, the capacity of the number of drives minus two can be used.

At each RAID level, read performance can be expected to improve by the number of drives, but write performance depends on the number of IO instructions (IOs) to the connected drive and the parity generation processing time. The table below summarizes IOs and maximum bandwidth issued by the controller at each RAID level. The actual performance value depends on the processing performance of the controller and the interface bandwidth of the storage drive.

 

| RAID level | | IOs to the backend [Note] | | Sequential read transfer performance | Sequential write transfer performance |
|---|---|---|---|---|---|
| | | Random read | Random write | | |
| RAID0 | Stripe | 1 | 1 | N x SR | N x SW |
| RAID1 | Mirror | 1 | 2 | N x SR | N x SW /2 |
| RAID10 | Mirror + Stripe | 1 | 2 | N x SR | N x SW /2 |
| RAID5 | Single Parity | 1 | 4 | N x SR | (N-1) x SW |
| RAID6 | Double Parity | 1 | 6 | N x SR | (N-2) x SW |

N: Number of drives, SR: Single read performance, SW: Single write performance

[Note] The number of IO instructions executed to the drive connected to the RAID controller.

Random read performs a single read access from the drive recording the target data. RAID1/10 random write performs two write access to the data stripe for mirroring.

RAID5 random write performs a total of four access to the drive by the following processing.

1) Reading the original data stripe.
2) Reading the original parity stripe, and computing a new parity stripe from the original data stripe and write data.
3) Updating the data stripe.
4) Writing a new parity stripe.

 

# Measurement Details

## *Measurement Method*

Performance measurements of the PRIMERGY and PRIMEQUEST server storage subsystems are performed by modeling drive access in real-world application scenarios based on the following specification.

The specified items are as follows.

- Random access/sequential access ratio
- Read access/write access ratio
- Block size (KiB)
- Queue depth (the number of IO requests issued at one time)

The combination of specified values is called "load profile." The following five standard load profiles correspond to typical application scenarios.

| Standard load profile | Access | Access type | | Block size [KiB] | Application |
|---|---|---|---|---|---|
| | | Read | Write | | |
| File copy | Random | 50% | 50% | 64 | File copy |
| File server | Random | 67% | 33% | 64 | File server |
| Database | Random | 67% | 33% | 8 | Database (data transfer) Mail server |
| Streaming | Sequential | 100% | 0% | 64 | Database (log file), data backup, video streaming (partial) |
| Restore | Sequential | 0% | 100% | 64 | File restore |

The measurement in this document uses a tool called Iometer to measure the maximum value by changing the queue depth (the number of IO requests issued at one time) with the access pattern based on these standard load profiles.

In order to measure including the response time of the storage system, it is executed with "# of Worker Threads = 1".

The main measurement results are as follows.

- Throughput [MiB/s] Data transfer amount per second (in megabytes)
- Transaction [IO/s] IO processing per second
- Latency [ms] Average response time (in milliseconds)

"Data Throughput" value was used for load profiles with sequential access pattern, while "Transaction Rate" value was used for load profiles with random access pattern. Throughput and transaction are in direct proportion to each other and can be calculated mutually using the following formula.

| | |
|---|---|
| *Data throughput [MiB/s]* | *= Transaction rate [IO/s] x Block size [MiB]* |
| *Transaction rate [IO/s]* | *= Data throughput [MiB/s] / Block size [MiB]* |

 

## *Measurement Environment*

The measurements were made using the following hardware and software components.

| **SUT (System Under Test)** | | | |
|---|---|---|---|
| Hardware | | | |
| Model | | PRIMERGY RX2540 M7<br>PRIMERGY RX2540 M6<br>PRIMERGY RX2540 M5 | |
| Onboard | RX2540 M6 | Intel VROC SATA RAID Controller<br>Intel VROC NVMe RAID Controller | |
| | RX2540 M5 | Embedded MegaRAID [Intel C620 PCH SATA Controller]<br>Microsoft storage pool [Intel PCIe Device Adapter] | |
| RAID controller | | PSAS CP2100-8i<br>PSAS CP2200-16i<br>PRAID CP400i<br>PRAID CP500i<br>PRAID CP600i<br>PRAID EP400i<br>PRAID EP420i<br>PRAID EP540i<br>PRAID EP640i<br>PRAID EP680i<br>PRAID EP3254-8i<br>PRAID EP3258-16i | |
| Storage drive | SATA | Seagate ST1000NX0313<br>Micron MTFDDAK240TDS | 7.2krpm, 512e<br>Read Intensive |
| | SAS | Toshiba AL15SEB030N<br>Toshiba AL15SEB090N<br>Seagate XS400ME70084<br>Seagate XS400ME70094 | 10krpm, 512n<br>10krpm, 512n<br>Write Intensive<br>Write Intensive, SED |
| | NVMe | Intel SSDPE2KE032T8 | Mixed Use |
| **Software** | | | |
| Operating system | | Microsoft Windows Server 2019 Standard | |
| Management software | | ServerView RAID Manager 7.6.3 | |
| Measurement tool | | Iometer 1.1.0 | |

 

# Intel VROC SATA RAID Controller RX2540 M6 onboard

SATA drives connected to the onboard SATA interface can be used with RAID 0/1/10. Up to 8 SATA drives can be connected.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| RAID0: 128 KiB<br>RAID1: N/A<br>RAID10: 64 KiB | Write Cache | off | Disabled |
| | Cache Policy | off | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #1 | Full size | RAW | Sequential test volume<br>Random test volume |

## *SATA-HDD Connection*

With RAID0, the performance value is proportional to the number of connected drives. With RAID1/10, read performance tends to be slower than the expected performance, but the write performance reached the expected performance.

[RAID0]

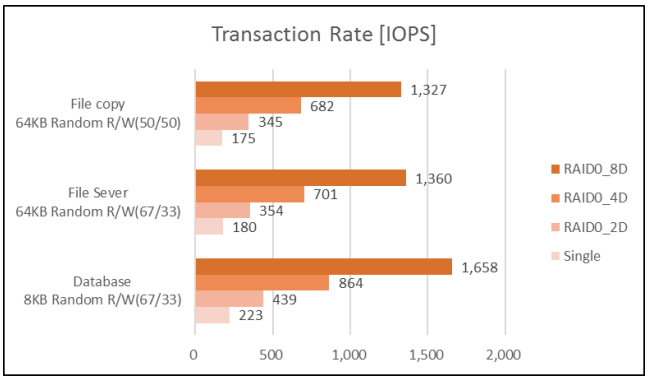


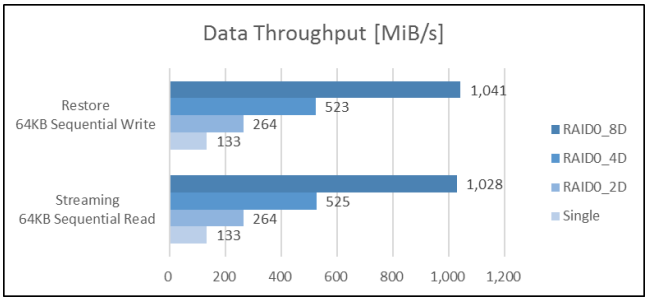


 

[RAID1/10]

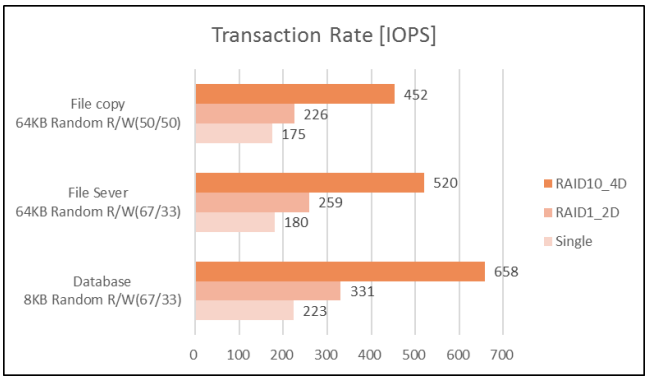


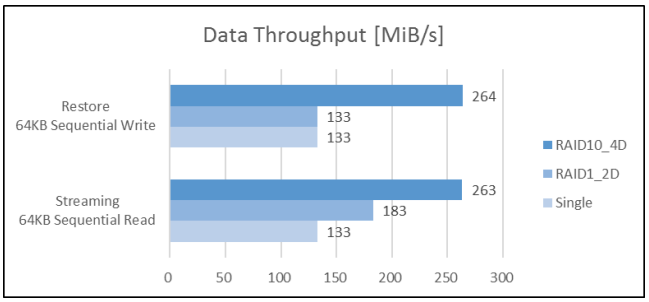


## *SATA-SSD Connection*

With RAID0, the performance value is proportional to the number of connected drives. RAID1/10 reached almost the expected performance.

[RAID0]

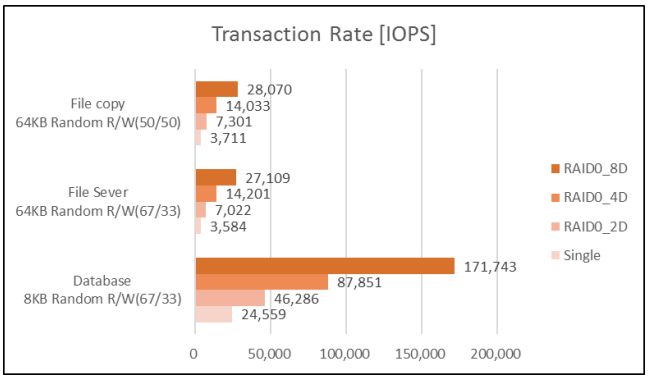


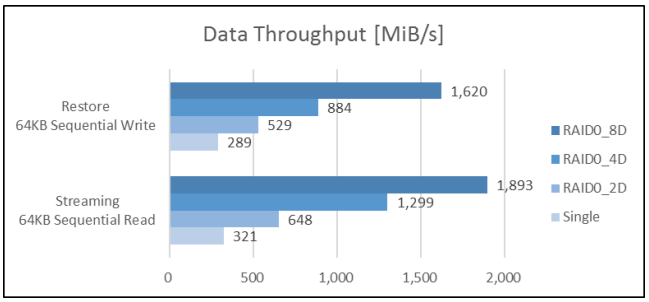


[RAID1/10]

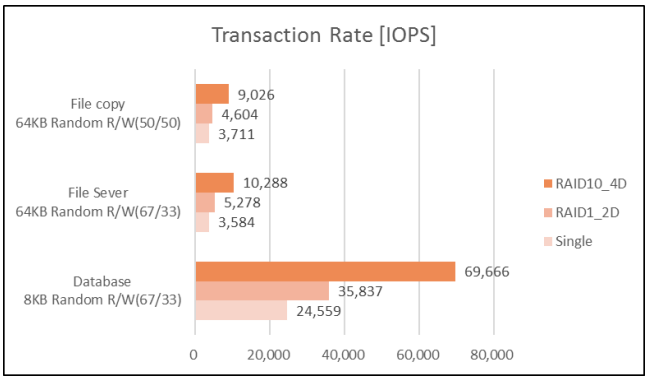


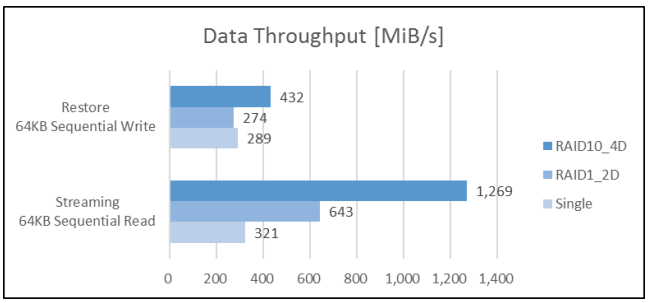


 

# Intel VROC NVMe RAID Controller RX2540 M6 onboard

NVMe drives connected to the PCIe lanes of Intel Xeon Scalable Processors can be used with RAID 0/1/10/5. Up to eight NVMe drives can be connected with two processors.

You must have the “Intel VROC Upgrade Key Premium” to use the RAID feature.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| RAID0: 128 KiB<br>RAID1: N/A<br>RAID10: 64 KiB<br>RAID5 x3: 64 KiB<br>RAID5 x4: 32 KiB<br>RADI5 x8: 16 KiB | Write Cache | off | Disabled |
| | Cache Policy | off | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #1 | Full size | RAW | Sequential test volume<br>Random test volume |

## *NVMe Connection*

Sequential performance has reached the limit when 4 drives are connected and is almost the same as when 8 drives are connected.

Since the performance value of the database pattern, which is a random access, is measured with "# of Worker Threads = 1", the latency of the transfer process becomes a bottleneck, and the transaction rate has peaked.

RAID 1/10 performance is nearly as expected with sequential access.

[RAID0]

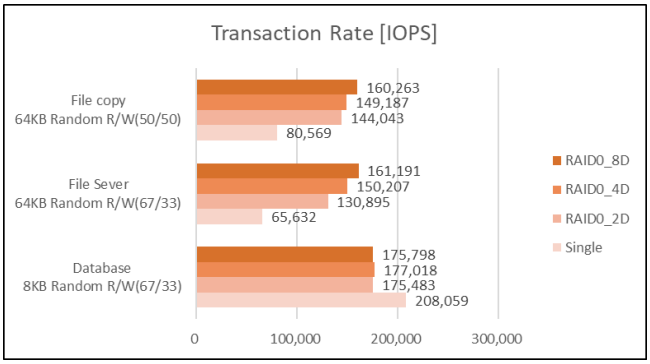


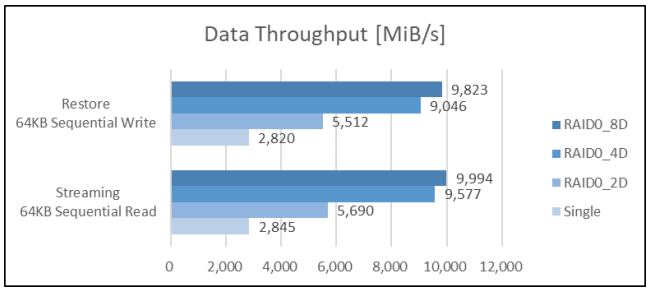


 

## [RAID1/10]

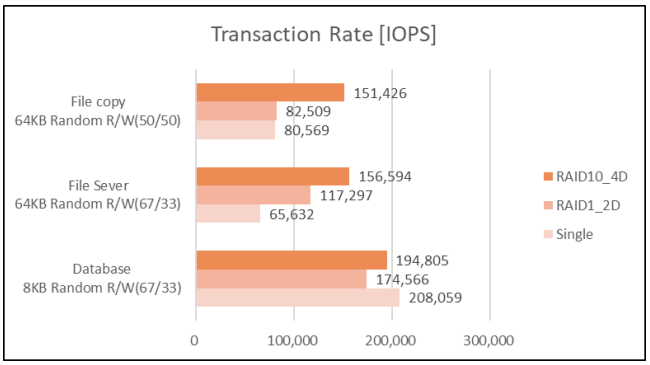


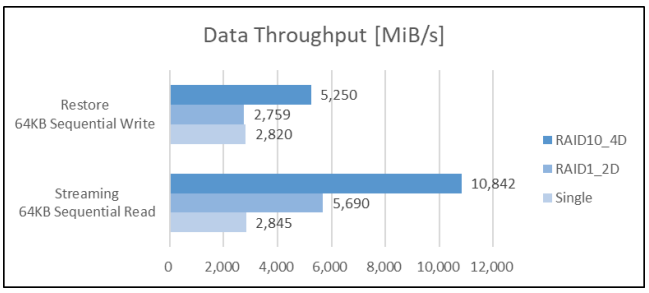


## [RAID5]

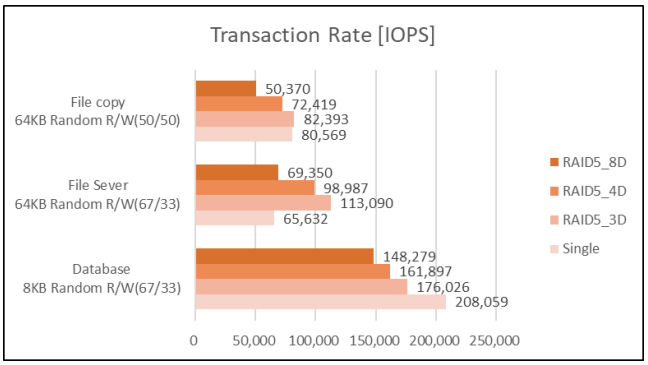


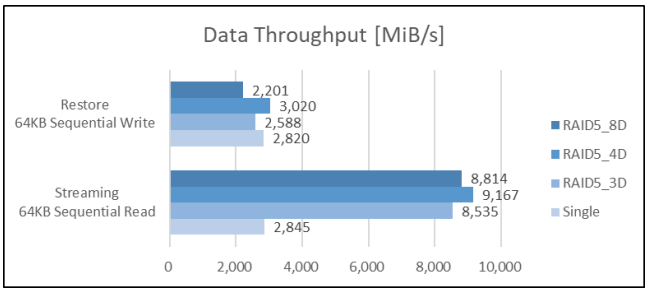


 

# Embedded MegaRAID [Intel C620 PCH SATA Controller] RX2540 M5 onboard

SATA drives connected to the onboard SATA interface can be used with RAID 0/1/10. Up to 8 SATA drives can be connected.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| 64 KiB | Read | Read Ahead | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #1 | full | RAW | Sequential test volume<br>Random test volume |

## *SATA-HDD Connection*

RAID0 read access has a performance value proportional to the number of connected drives, but write performance remains a low rate of increase. RAID1 is almost the same as single drive performance, and with RAID10, the write performance remains at the same low rate of increase as RAID0.

[RAID0]

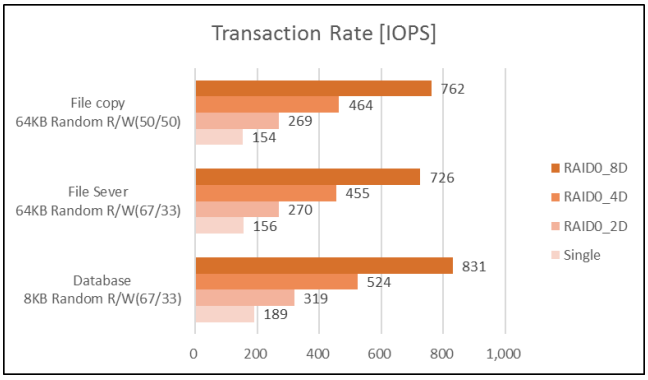


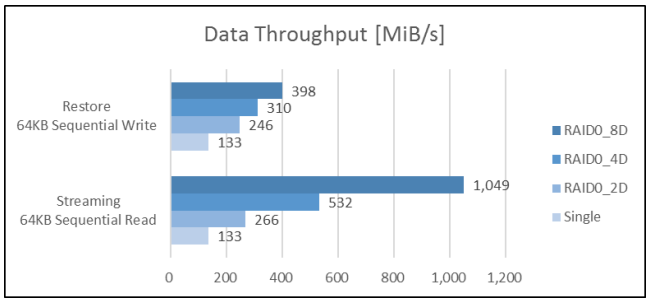


 

[RAID1/10]

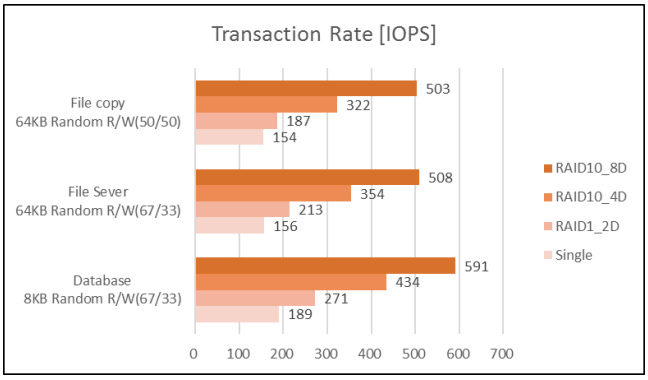


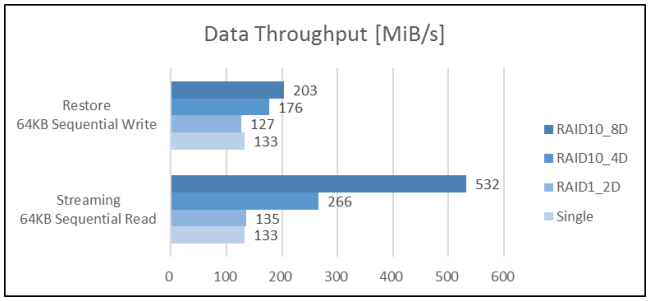


## *SATA-SSD Connection*

With RAID0, the performance value is proportional to the number of connected drives. RAID1/10 is almost the same as the expected performance.

[RAID0]

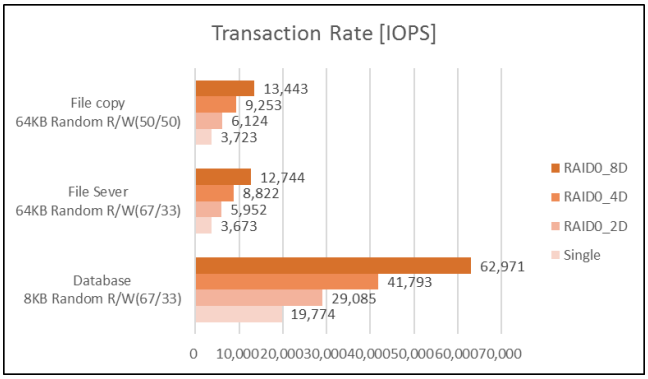


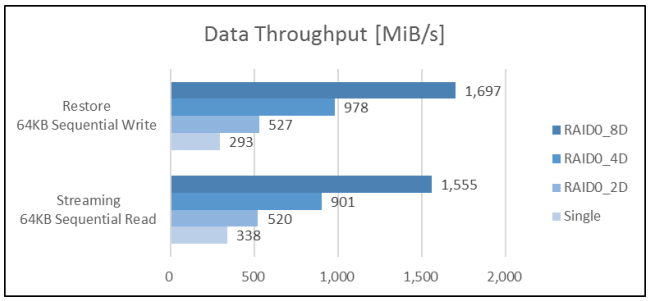


[RAID1/10]

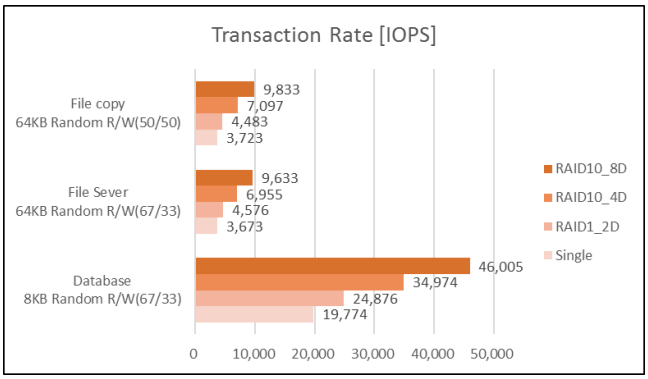


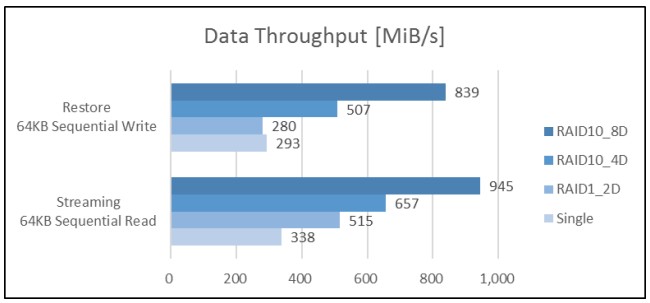


 

# Microsoft storage pool [Intel PCIe Device Adapter] RX2540 M5 onboard

NVMe drives connected to the PCIe interface of the Intel Xeon scalable Processor can be used as virtual drives configured by the Microsoft storage pool. You can specify Simple mode (equivalent to RAID0) or Mirror mode (equivalent to RAID1/10). Up to 8 NVMe drives can be connected by two processors.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| 256 KiB | Read | - | Enabled [Note] |
| | Write | Enabled | |

[Note] The cache is enabled, but used NVMe drives do not have a write cache.

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #1 | 64 GiB | NTFS | Sequential test volume<br>Random test volume |

## *NVMe Connection*

Sequential performance has reached the limit when 4 drives are connected and is almost the same as when 8 drives are connected.

Since the performance value of the database pattern, which is a random access, is measured with "# of Worker Threads = 1", the latency of the transfer process becomes a bottleneck, and the transaction rate has peaked.

RAID1 has the same performance as a single drive. RAID10 has more than double the read performance of RAID1, but writes tend to be low.

[Simple]

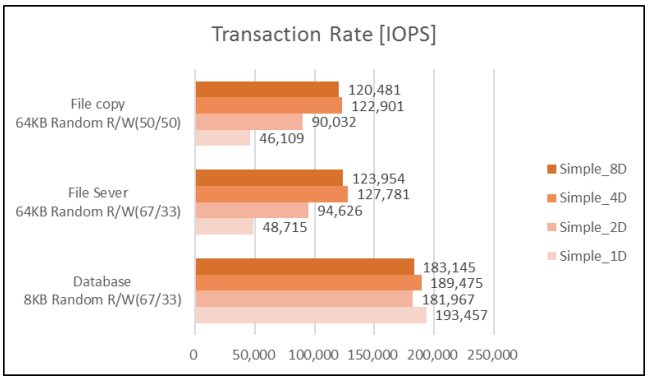


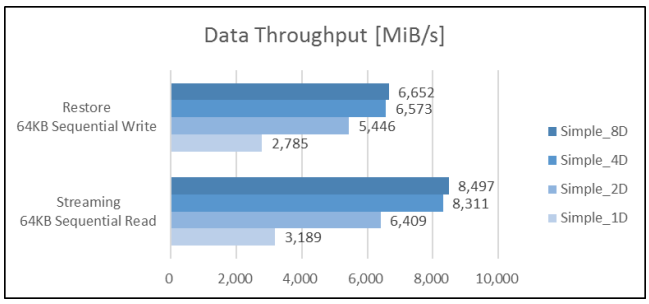


 

## [Mirror]

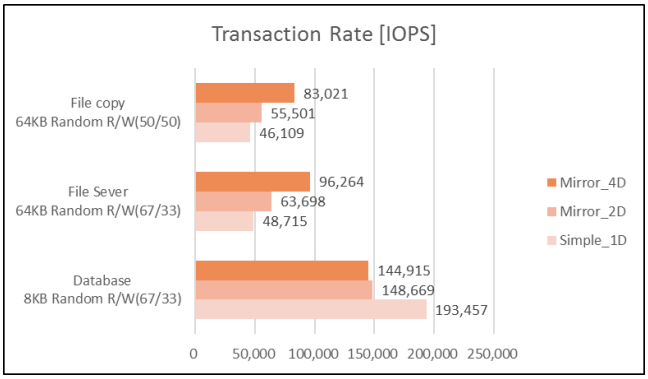


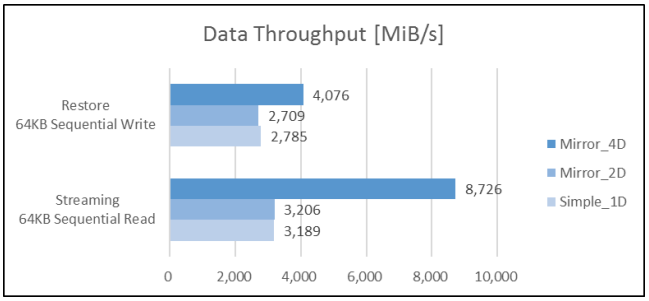


 

# PSAS CP2100-8i
# PCIe option

Up to 8 SAS/SATA drives can be directly connected. The number of drives that can be connected depends on the server configuration. The measurements were performed by connecting 16 drives using a SAS switch.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | Disk Cache |
|---|---|---|
| 256 KiB | - | Disabled |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #1 | 64 GiB | RAW | Sequential test volume<br>Random test volume |

## *SAS-HDD Connection (Measuring Up to 16 Drives via SAS Switch)*

With RAID0/1/10, the performance value is proportional to the number of connected drives.

With RAID5, read performance is proportional to the number of drives, but write performance does not improve.

[RAID0]

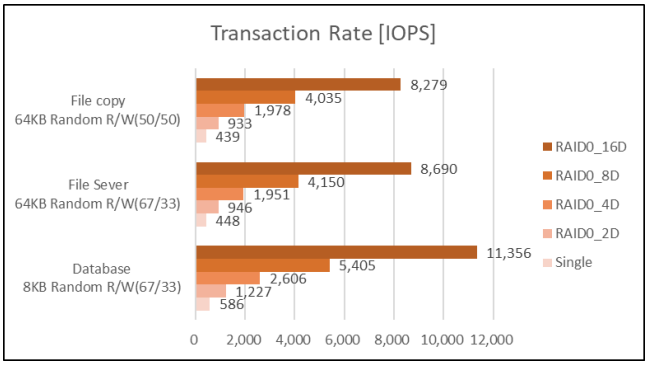


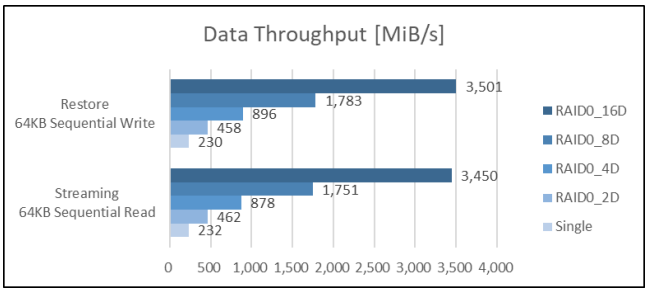


[RAID1/10]

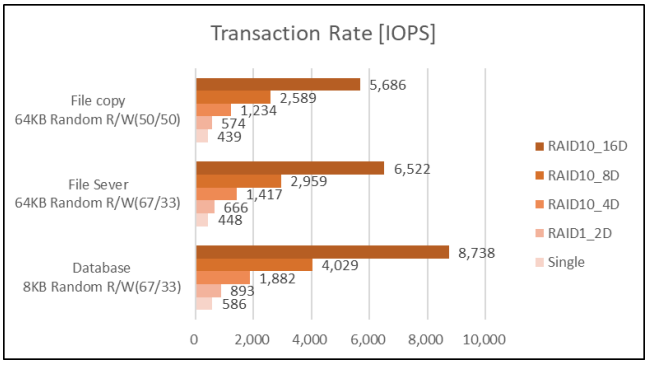


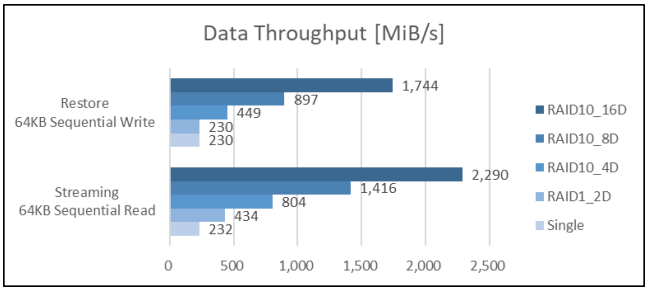


 

[RAID5]

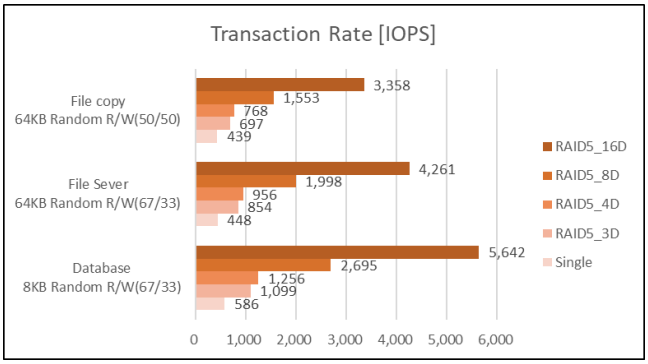


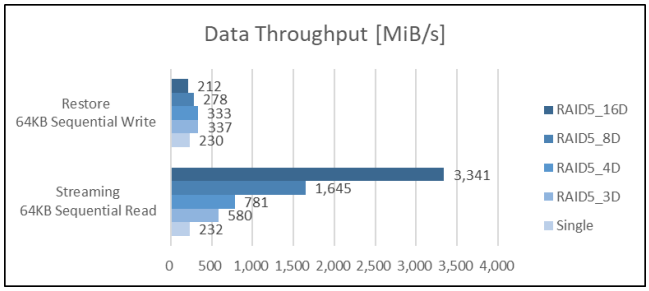


## *SAS-SSD Connection (Measuring Up to 16 Drives via SAS Switch)*

The sequential access performance of 16 drives is almost the same as that of 8 drives. By connecting with 8 drives of about 1,000 MiB/s per drive, the transfer capacity on the drive side can reach about 8,000 MiB/s. However, the theoretical throughput of PCIe Gen3 x8 lane, which is the interface on the CPU side, is 7,512 MiB/s, so it is presumed that the performance limit of the CPU side interface has been reached when eight drives are connected.

[RAID0]

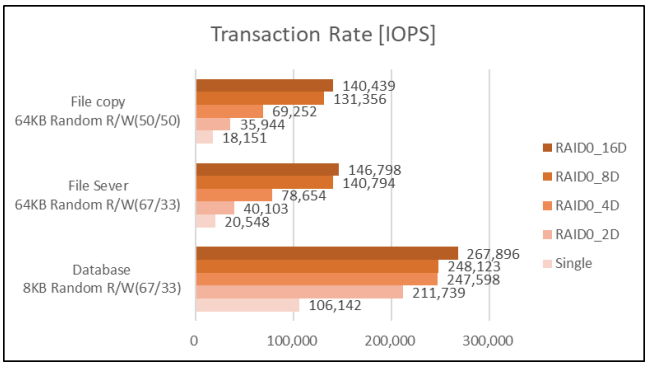


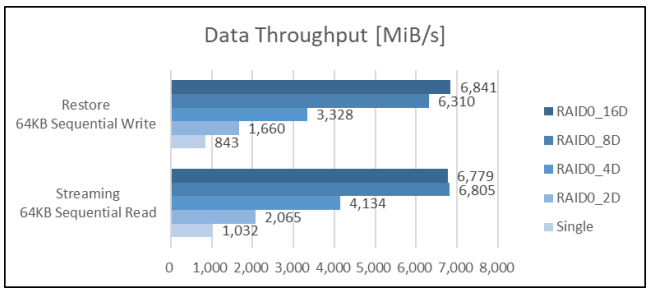


[RAID1/10]

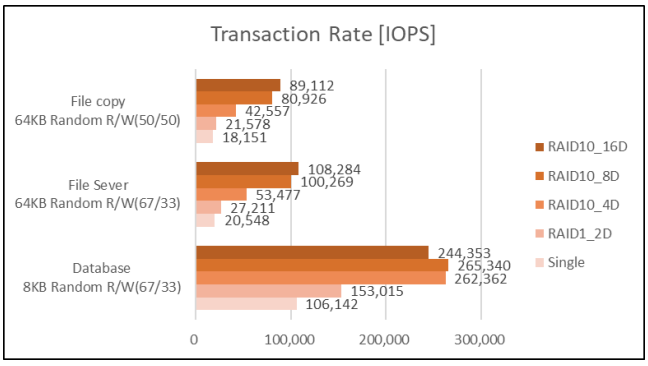


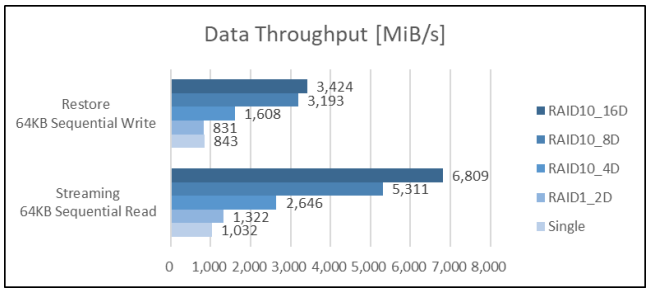


 

## [RAID5]

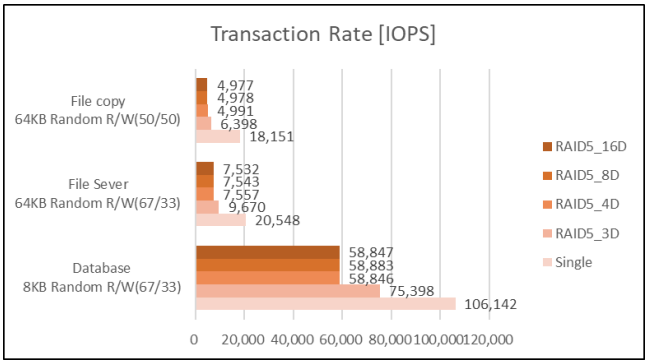


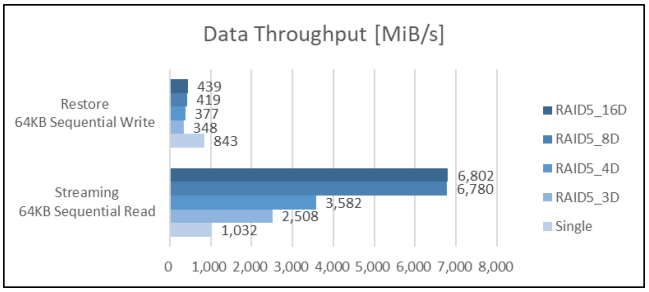


 

# PSAS CP2200-16i
# PCIe option

The host bus supports PCIe 4.0, and the data bus bandwidth is improved. Up to 16 SAS/SATA drives or 4 NVMe drives can be directly connected to the RAID controller. The number of drives that can be connected depends on the server configuration.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | Disk Cache |
|---|---|---|
| HDD: 256 KiB<br>SSD: 64 KiB | - | Disabled |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #1 | 64 GiB | RAW | Sequential test volume<br>Random test volume |

## *SAS-HDD Connection*

RAID0 performance values increases to the number of drives connected.

RAID1/10 sequential write performance is expected with 2/4/8 drives (single drive performance x number of drives/2 times), while 16 drives are almost as good as 8 drives.

RAID5 sequential write performance does not improve with more drives. This is presumably because the RAID controller's parity generation process is rate-limiting.

[RAID0]

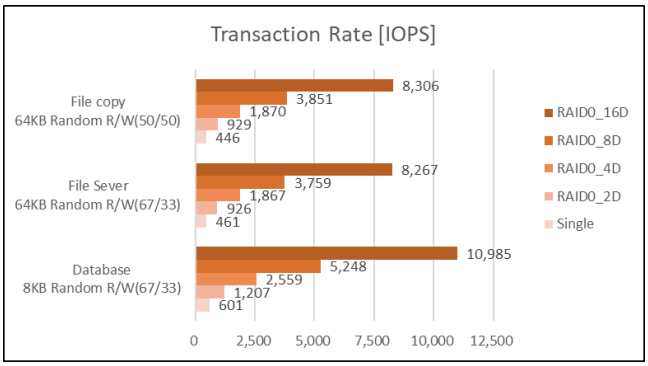


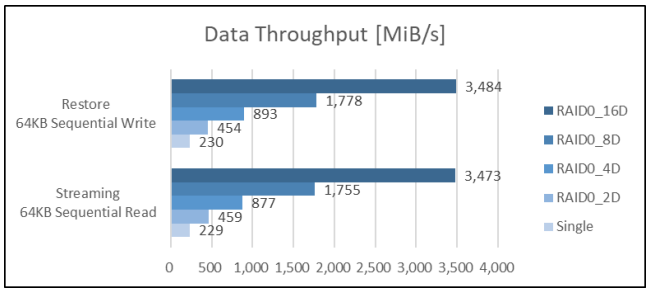


 

[RAID1/10]

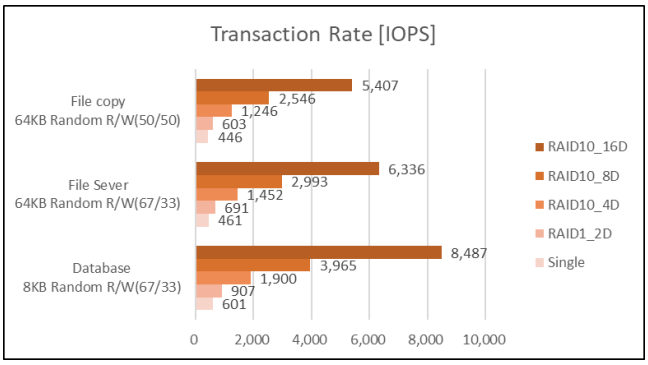


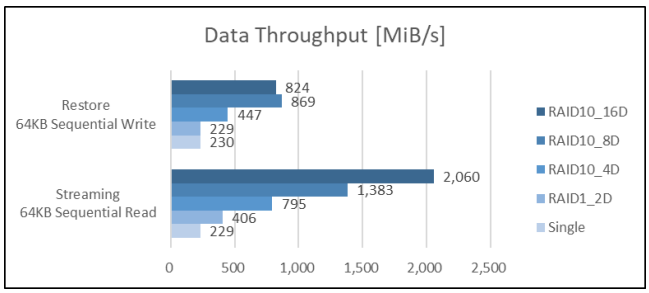


[RAID5]

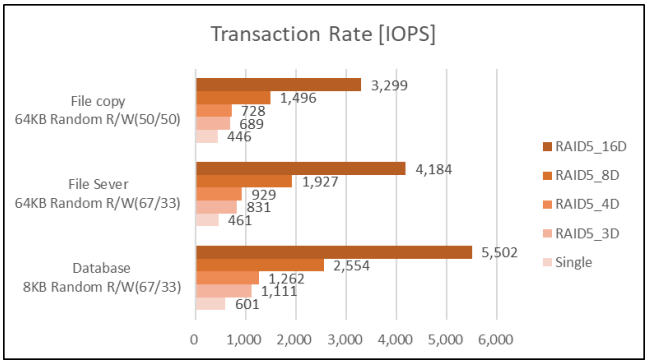


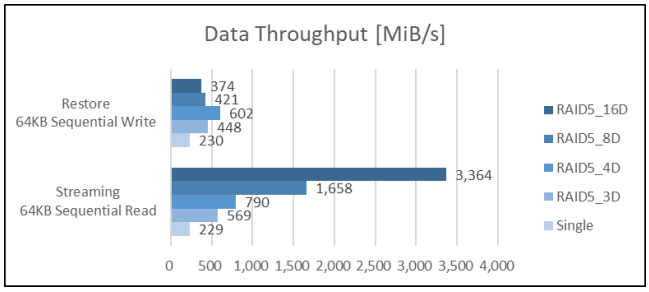


## *SAS-SSD Connection*

Increased bandwidth with PCIe 4.0 support on the host bus provides 10,000~12,000 MiB/s read performance with 16 drives. Since the performance value of the database pattern is measured with "# of Worker Threads=1" the latency of transfer processing becomes a bottleneck and the transaction rate reaches a plateau. With RAID5, the read performance is proportional to the number of drives, but the access pattern that contains writes is degraded compared to without RAID. This is presumably because the RAID controller's parity generation process is rate-limiting.

[RAID0]

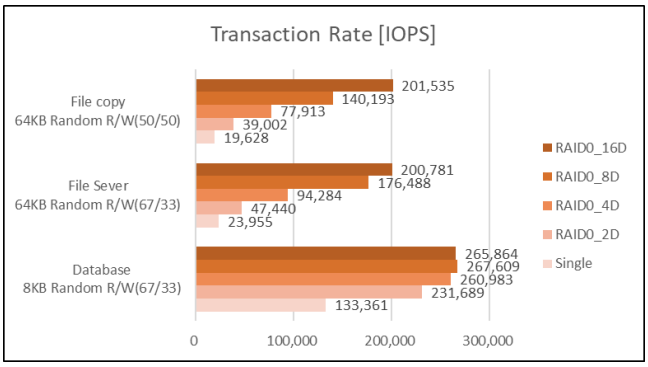


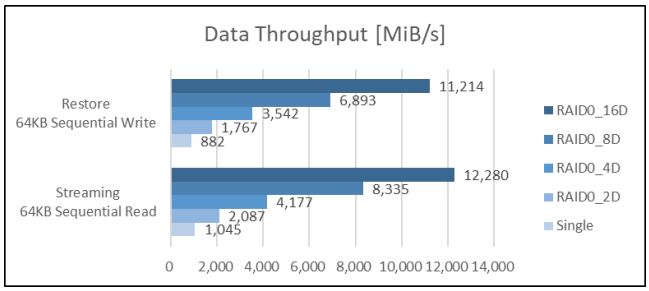


 

[RAID1/10]

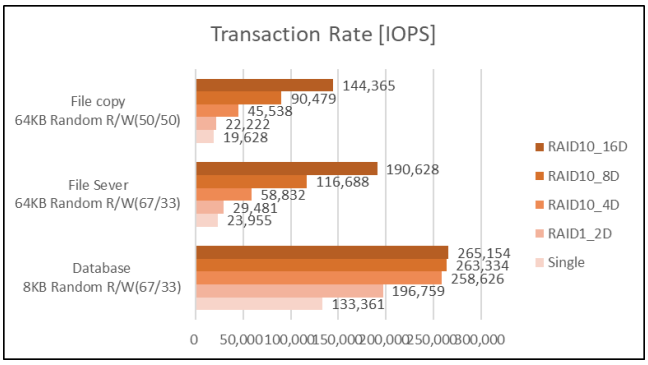


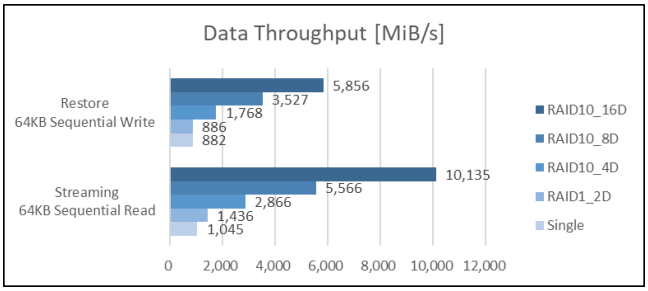


[RAID5]

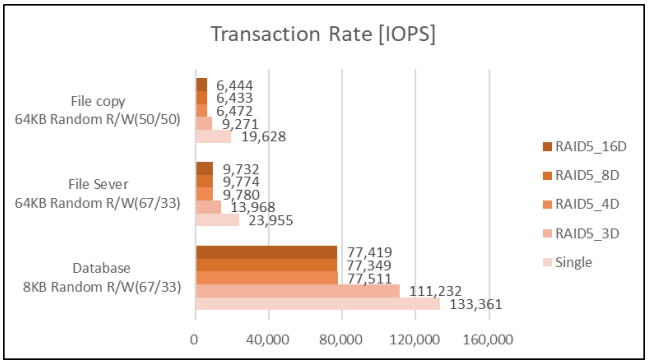


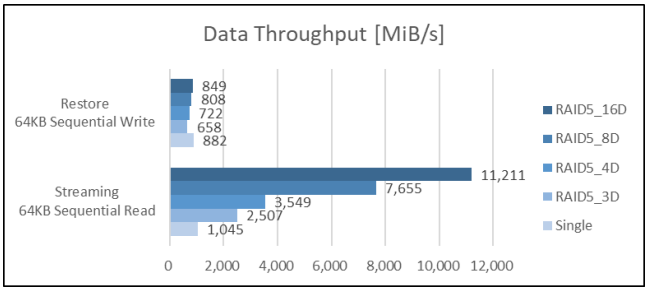


## *NVMe Connection*

Increased bandwidth with PCIe 4.0 support on the host bus provides 10,000~12,000 MiB/s read performance with 16 drives. Since the performance value of the database pattern is measured with "# of Worker Threads=1" the latency of transfer processing becomes a bottleneck and the transaction rate reaches a plateau. With RAID5, the read performance is proportional to the number of drives, but the access pattern that contains writes is degraded compared to without RAID. This is presumably because the RAID controller's parity generation process is rate-limiting.

[RAID0]

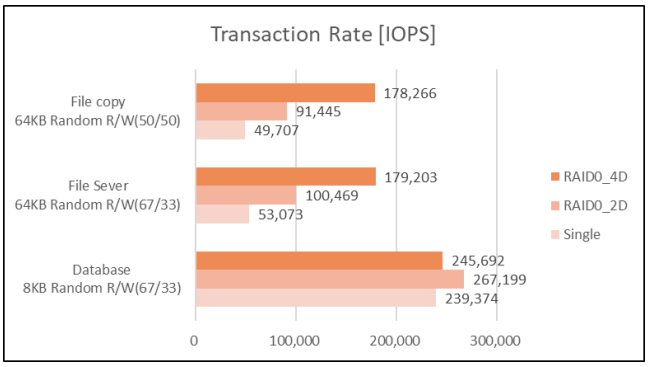


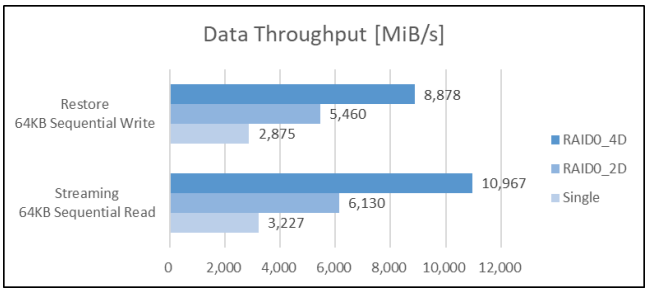


 

## [RAID1/10]

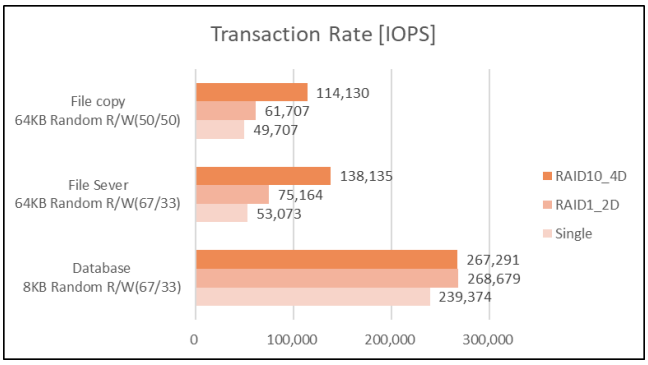


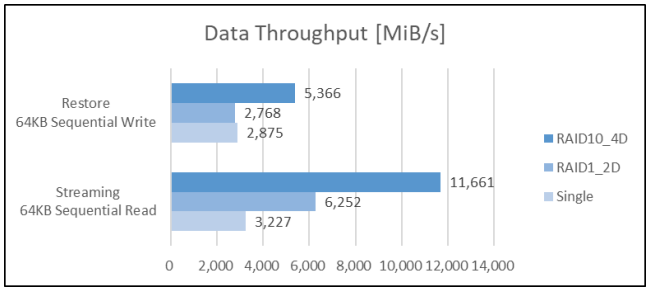


## [RAID5]

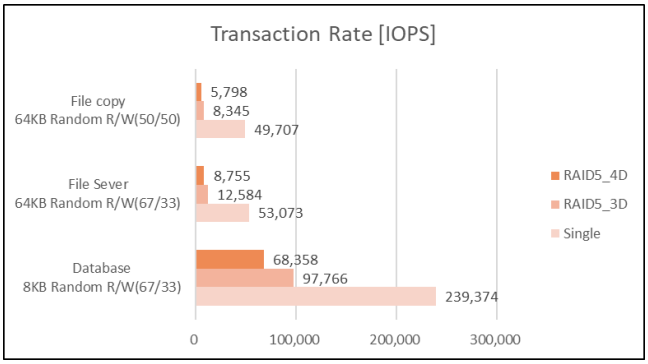


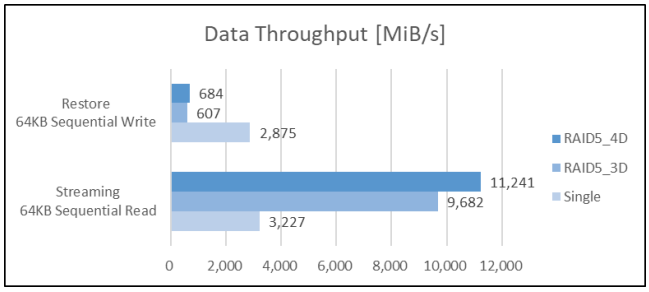


 

# PRAID CP400i
# PCIe option

Up to 8 SAS/SATA drives can be directly connected. The number of drives that can be connected depends on the server configuration. The measurements were performed by connecting 16 drives using a SAS switch.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| 64 KiB | Read | No Read Ahead | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #2 | 32 GiB | RAW | Sequential test volume |
| | 64 GiB | RAW | Random test volume |

## *SAS-HDD Connection (Measuring Up to 16 Drives via SAS Switch)*

With RAID 0/RAID 5, performance is proportional to the number of drives connected, but over 8 drives, sequential write performance tend to plateau.

The sequential performance of RAID1 is almost the same as the performance of a single drive. RAID10 has a performance proportional to the performance of RAID1 (double for four drives and 4 times for eight drives).

[RAID0]

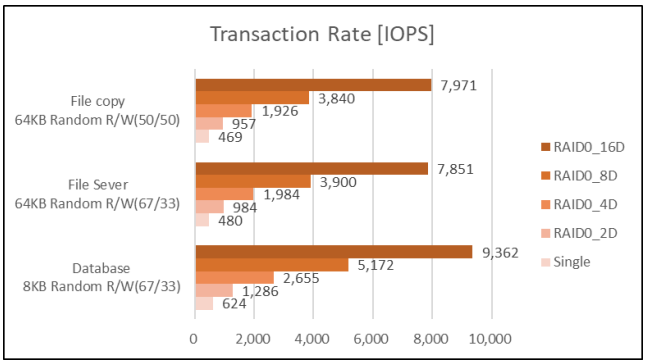


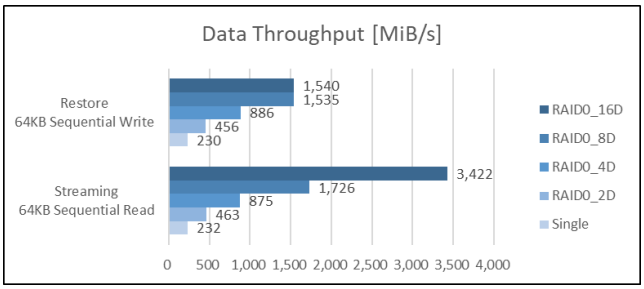


 

[RAID1/10]

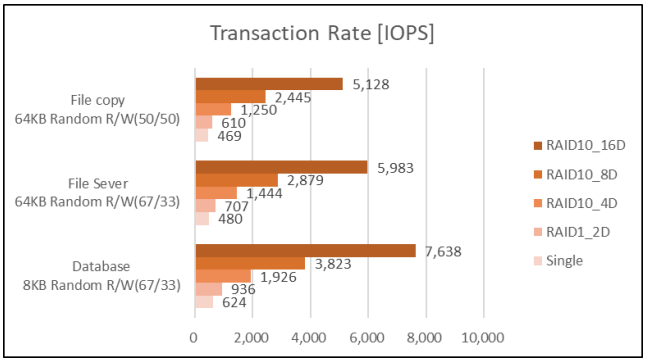


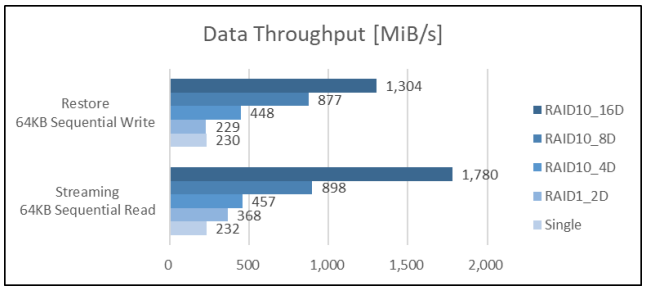


[RAID5]

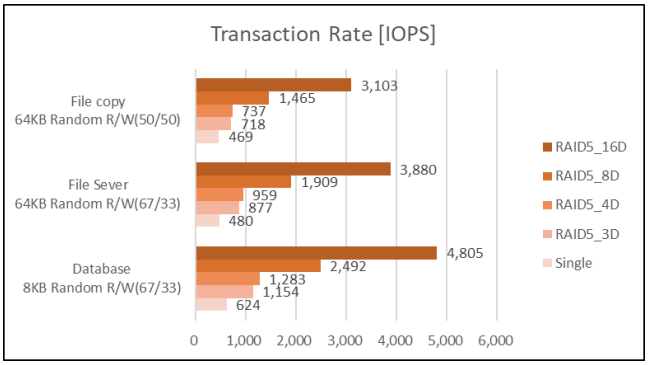


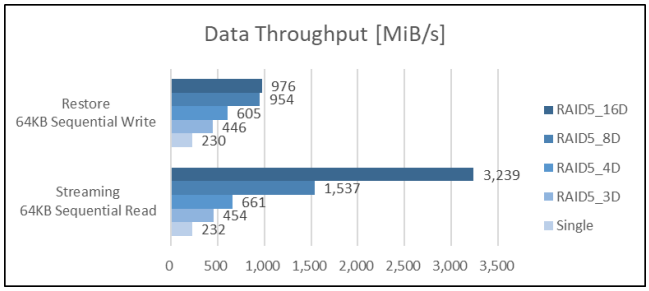


## *SAS-SSD Connection (Measuring Up to 16 Drives via SAS Switch)*

The sequential access performance of 16 drives is almost the same as that of 8 drives. By connecting with 8 drives of about 1,000 MiB/s per drive, the transfer capacity on the drive side can reach about 8,000 MiB/s. However, the theoretical throughput of PCIe Gen3 x8 lane, which is the interface on the CPU side, is 7,512 MiB/s, so it is presumed that the performance limit of the CPU side interface has been reached when eight drives are connected.

In addition, a throughput of about 3,200 MiB/s is obtained with RAID1/10 write performance, and a throughput of about 1,900 MiB/s is obtained with RAID5 write performance.

[RAID0]

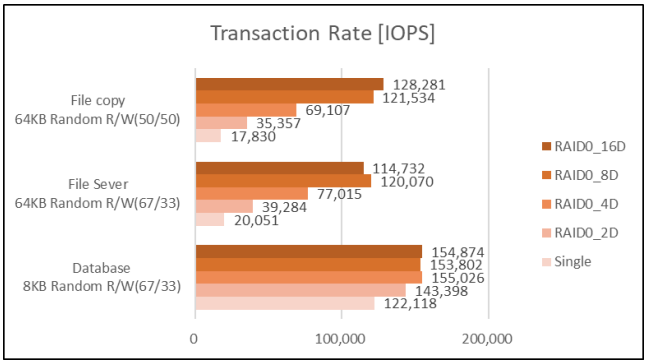


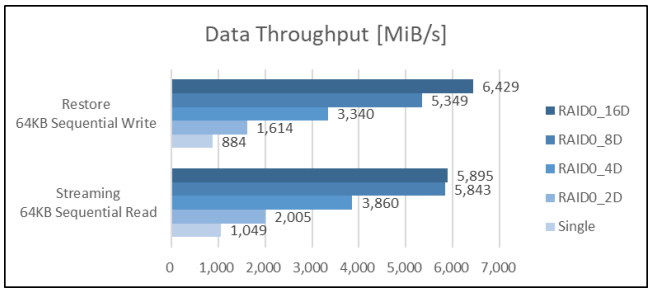


 

## [RAID1/10]

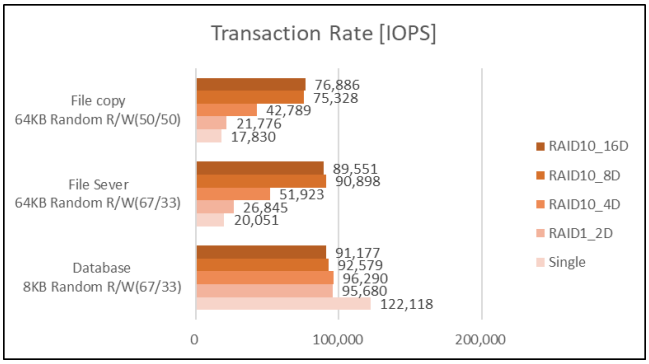


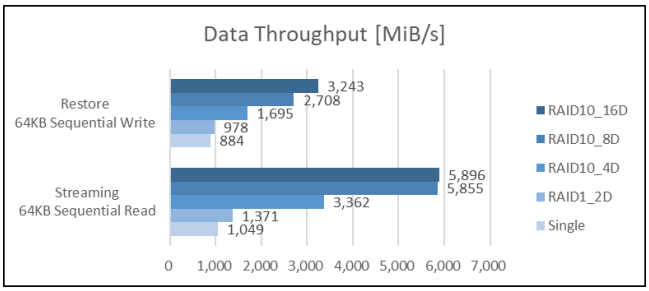


## [RAID5]

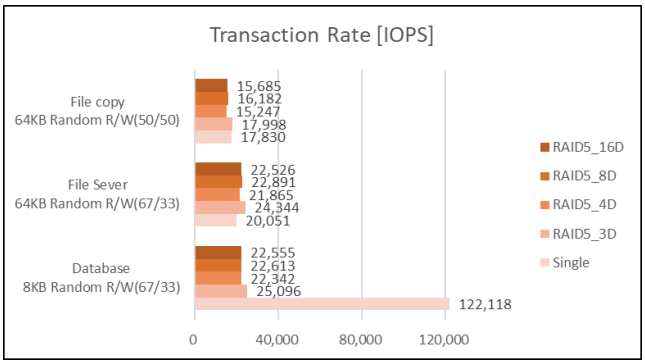


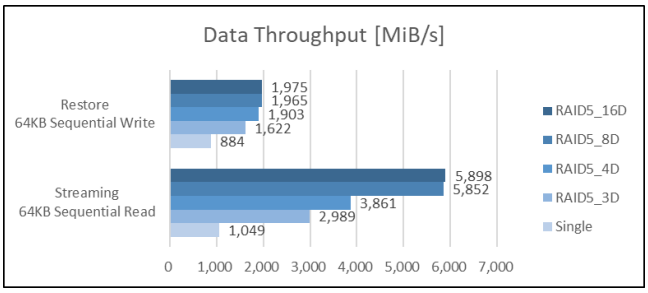


 

# PRAID CP500i
# PCIe option

Up to 8 SAS/SATA drives can be directly connected. The number of drives that can be connected depends on the server configuration. The measurements were performed by connecting 16 drives using a SAS switch.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| 64 KiB | Read | No Read Ahead | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #2 | 32 GiB | RAW | Sequential test volume |
| | 64 GiB | RAW | Random test volume |

## *SAS-HDD Connection (Measuring Up to 16 Drives via SAS Switch)*

With RAID0/RAID5, the performance value is proportional to the number of connected drives.

The sequential performance of RAID1 is almost the same as the performance of a single drive. RAID10 has a performance proportional to the performance of RAID1 (double for four drives and 4 times for eight drives).

[RAID0]

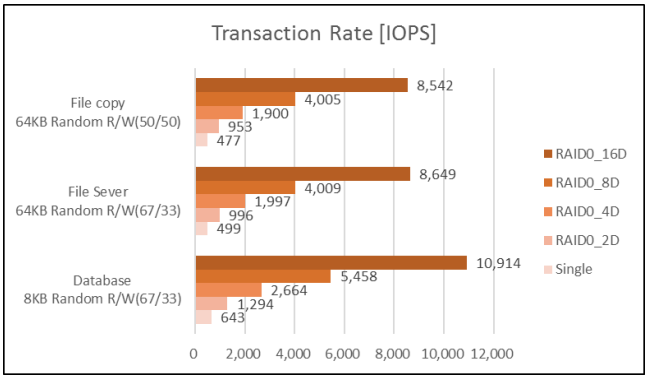


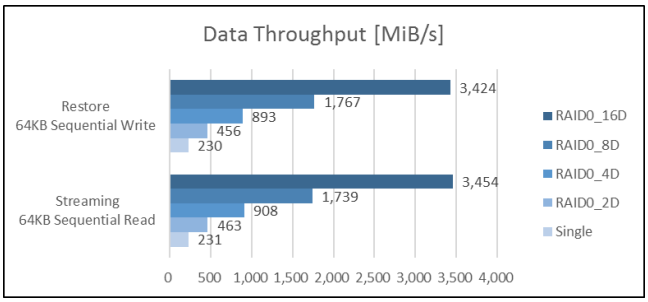


 

[RAID1/10]

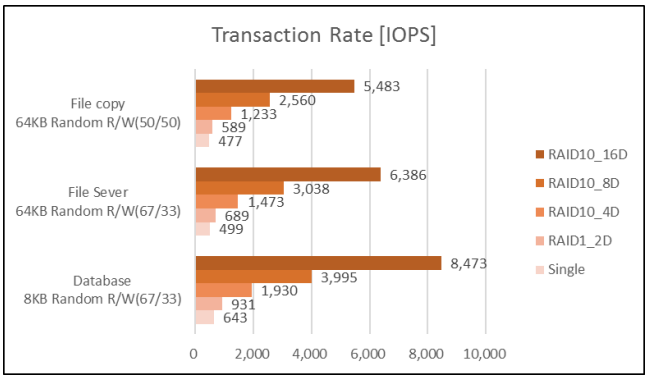


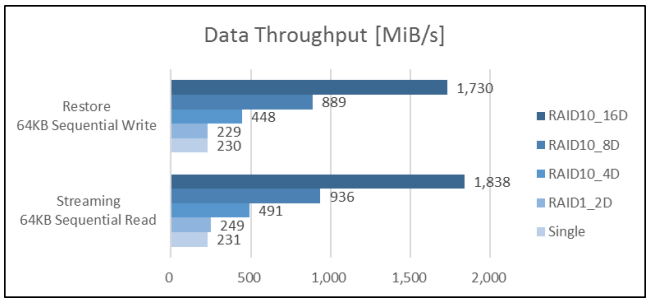


[RAID5]

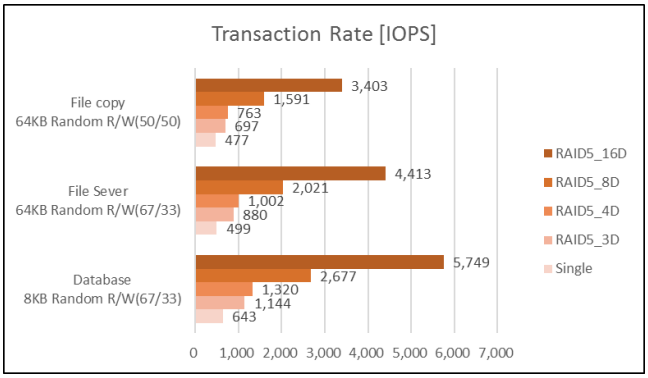


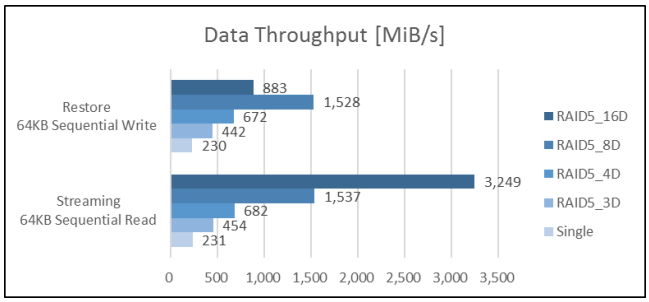


## *SAS-SSD Connection (Measuring Up to 16 Drives via SAS Switch)*

The sequential access performance of 16 drives is almost the same as that of 8 drives. By connecting with 8 drives of about 1,000 MiB/s per drive, the transfer capacity on the drive side can reach about 8,000 MiB/s. However, the theoretical throughput of PCIe Gen3 x8 lane, which is the interface on the CPU side, is 7,512 MiB/s, so it is presumed that the performance limit of the CPU side interface has been reached when eight drives are connected.

In addition, a throughput of about 2,300 MiB/s is obtained with RAID1/10 write performance, and a throughput of about 2,200 MiB/s is obtained with RAID5 write performance.

[RAID0]

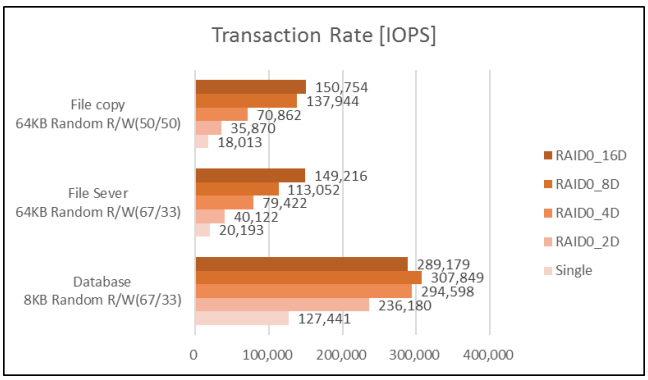


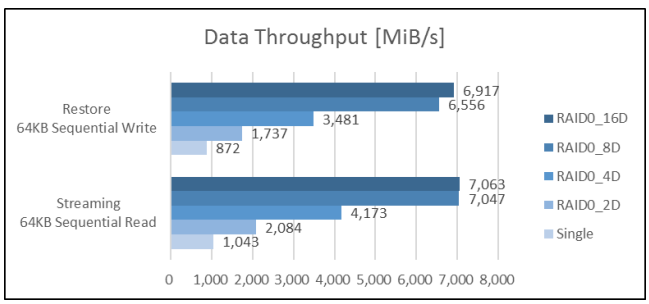


 

## [RAID1/10]

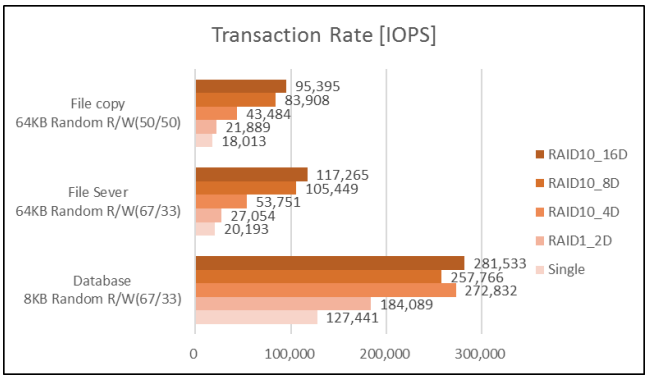


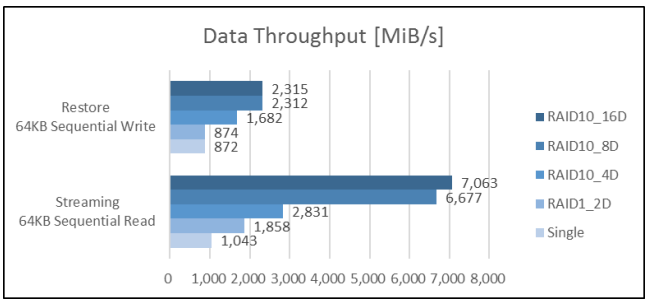


## [RAID5]

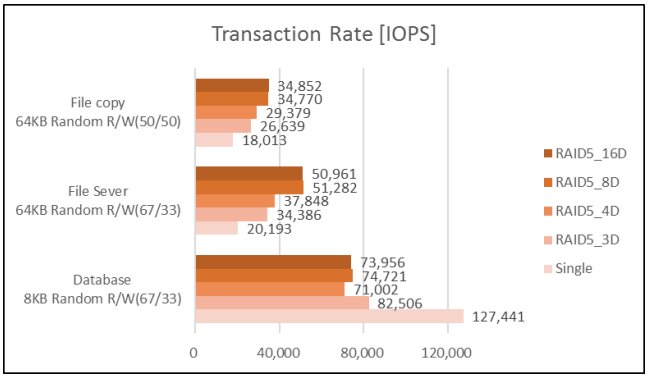


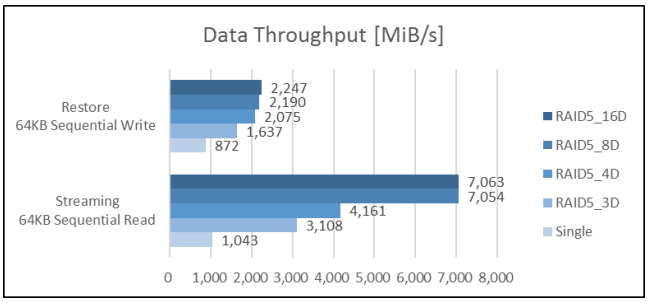


 

# PRAID CP600i
# PCIe option

The host bus supports PCIe 4.0, and the data bus bandwidth is improved. Up to 8 SAS/SATA drives can be directly connected. The number of drives that can be connected depends on the server configuration. The measurements were performed by connecting 16 drives using a SAS switch.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| 64 KiB | Read | No Read Ahead | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #2 | 32 GiB | RAW | Sequential test volume |
| | 64 GiB | RAW | Random test volume |

## *SAS-HDD Connection (Measuring Up to 16 Drives via SAS Switch)*

With RAID0, the performance value is proportional to the number of connected drives.

The sequential performance of RAID1 is almost the same as the performance of a single drive. RAID10 has a performance proportional to the performance of RAID1 (double for four drives and 4 times for eight drives).

[RAID0]

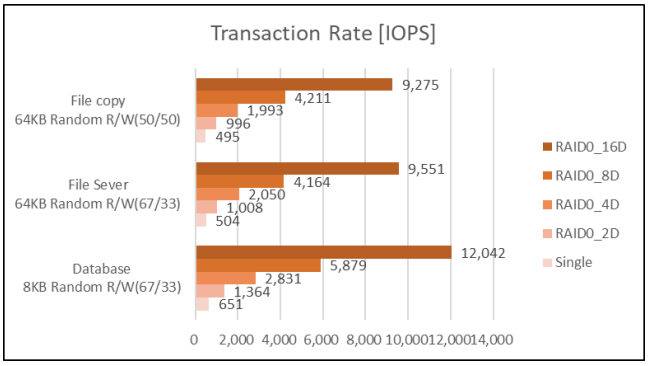


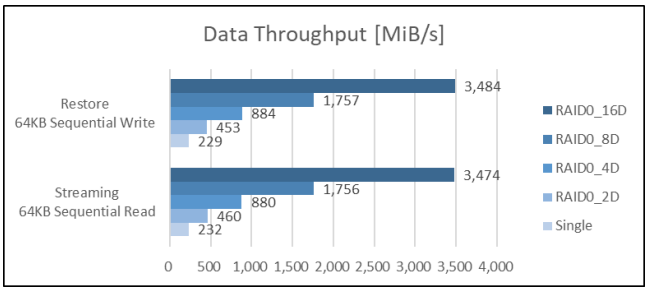


 

[RAID1/10]

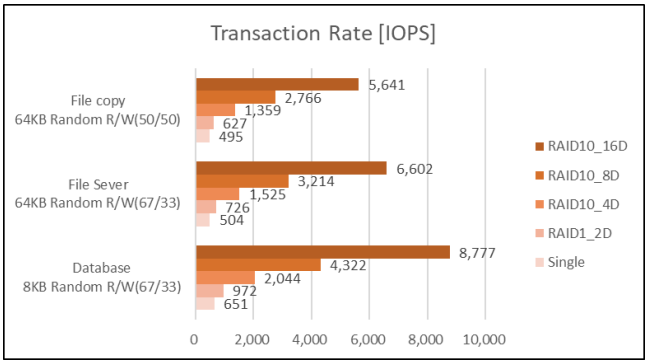


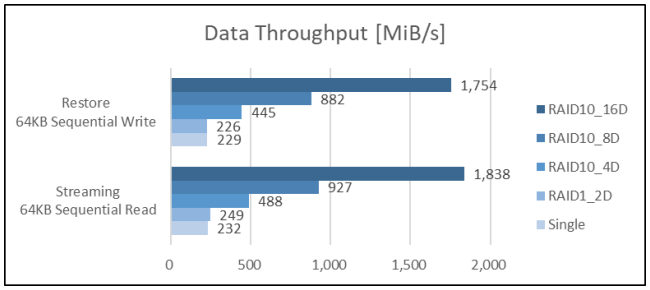


## *SAS-SSD Connection (Measuring Up to 16 Drives via SAS Switch)*

The sequential access performance of 16 drives is almost the same as that of 8 drives. Since the number of ports of drive-side is 8 ports, it is assumed that the drive-side throughput limit of about 8,000 MiB/s was reached at the time of 8 drives were connected. Since the performance limit has been reached, increasing the number of drives will not improve performance.

In addition, a throughput of about 4,000 MiB/s is obtained with RAID1/10 write performance.

[RAID0]

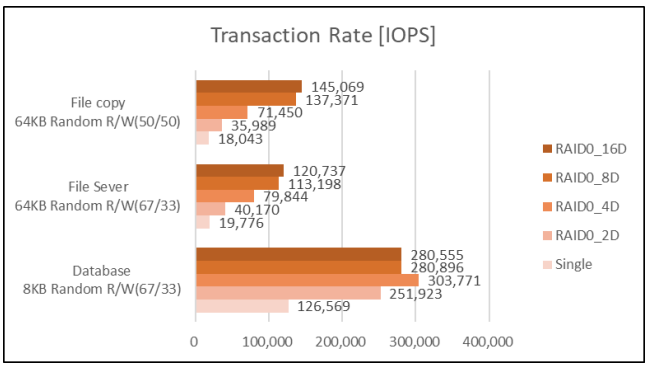


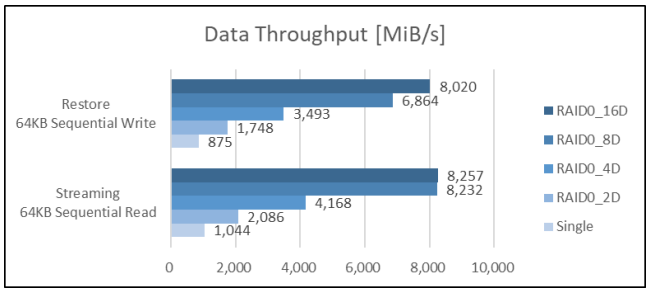


[RAID1/10]

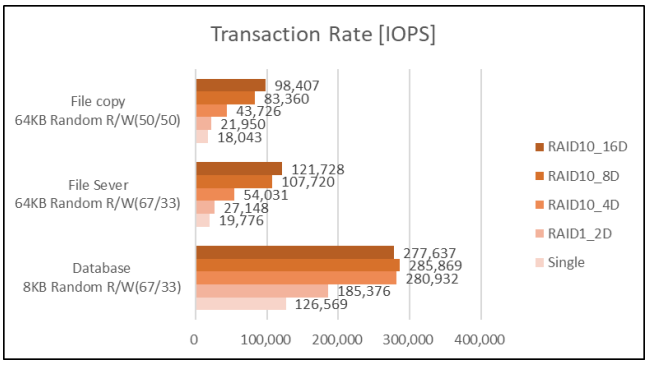


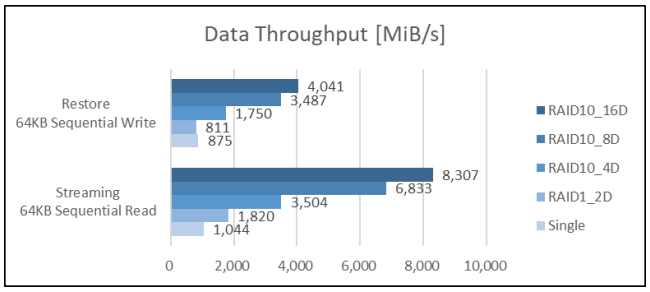


 

# PRAID EP400i
# PCIe option

Up to 8 SAS/SATA drives can be connected drives directly. The number of drives that can be connected depends on the server configuration. Buffering by cache memory can be expected to improve the performance of the write processing.

In this measurement, in order to see the basic characteristics of the RAID controller, the measurements were performed without using the buffering function by the cache memory.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| HDD: 256 KiB<br>SSD: 64 KiB | Read | No Read Ahead | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #2 | 32 GiB | RAW | Sequential test volume |
| | 64 GiB | RAW | Random test volume |

## *SAS-HDD Connection (Measuring Up to 16 Drives via SAS Switch)*

With RAID0/RAID5, the performance value is proportional to the number of connected drives.

Since the read performance of RAID1 is higher than that of a single drive, the read performance of RAID10 is also improved.

[RAID0]

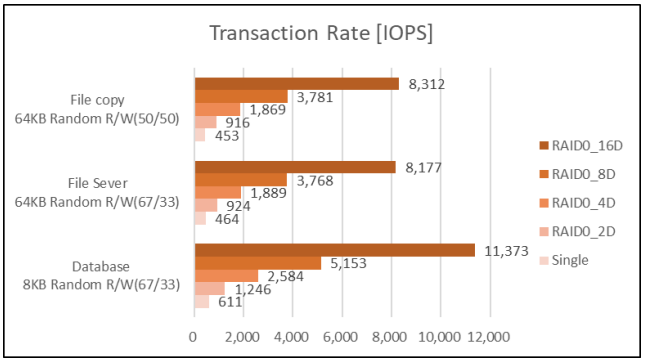


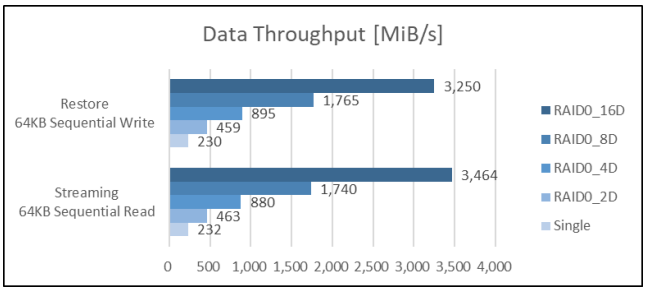


 

[RAID1/10]

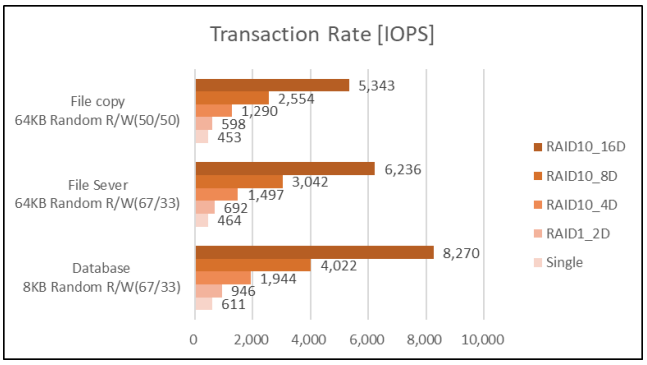


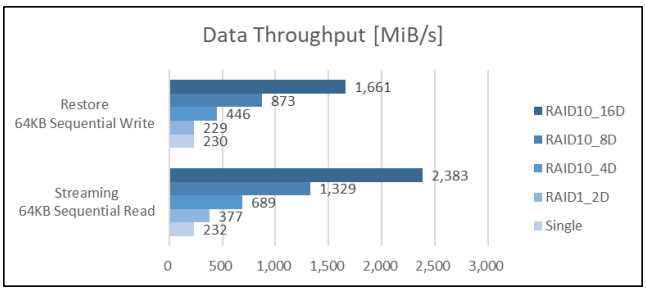


[RAID5]

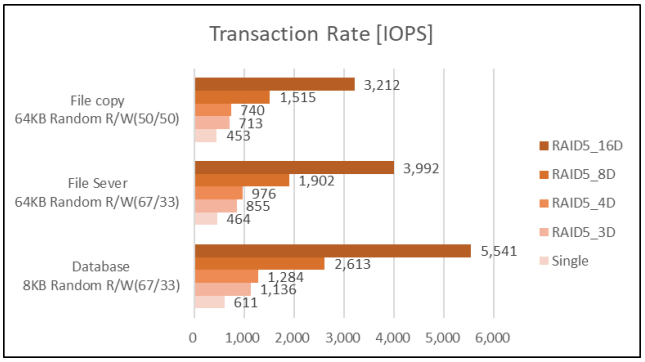


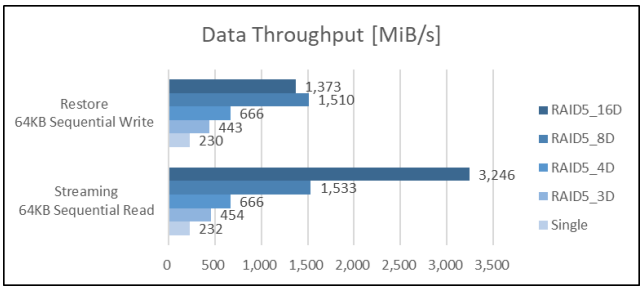


## *SAS-SSD Connection (Measuring Up to 16 Drives via SAS Switch)*

The sequential access performance of 16 drives is almost the same as that of 8 drives. By connecting with 8 drives of about 1,000 MiB/s per drive, the transfer capacity on the drive side can reach about 8,000 MiB/s. However, the theoretical throughput of PCIe Gen3 x8 lane, which is the interface on the CPU side, is 7,512 MiB/s, so it is presumed that the performance limit of the CPU side interface has been reached when eight drives are connected.

In addition, a throughput of about 3,200 MiB/s is obtained with RAID1/10 write performance, and a throughput of about 3,100 MiB/s is obtained with RAID5 write performance.

[RAID0]

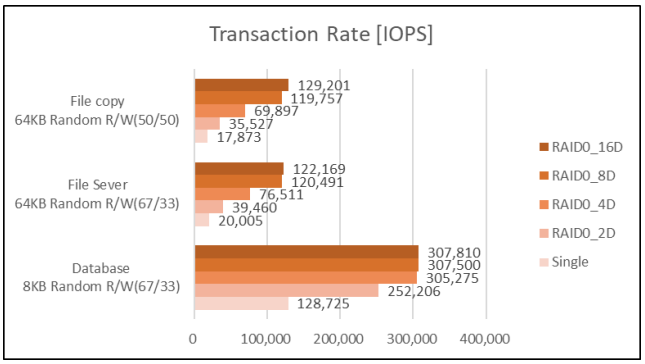


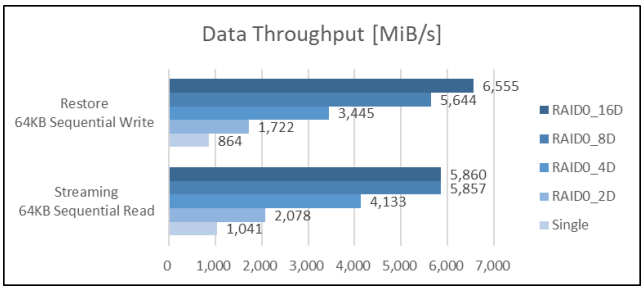


 

## [RAID1/10]

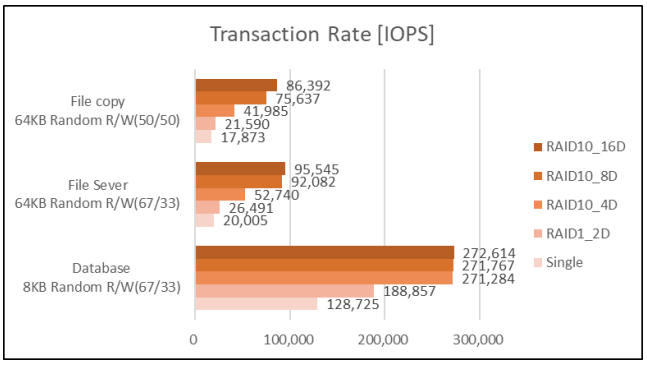


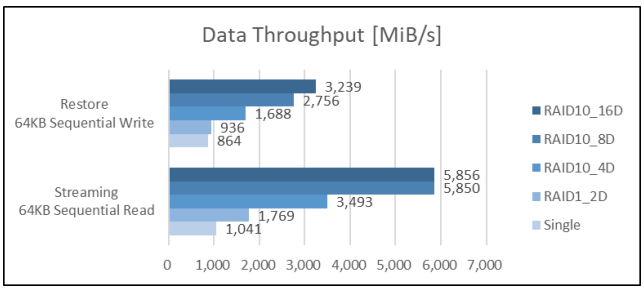


## [RAID5]

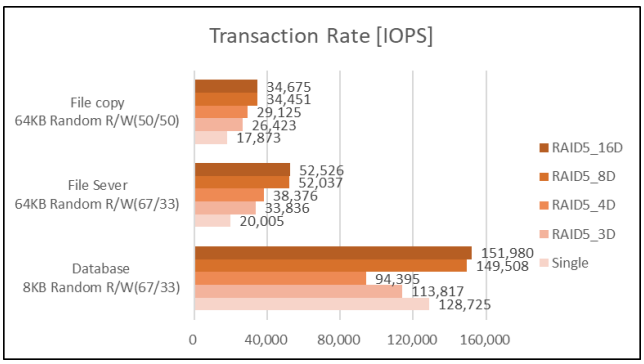


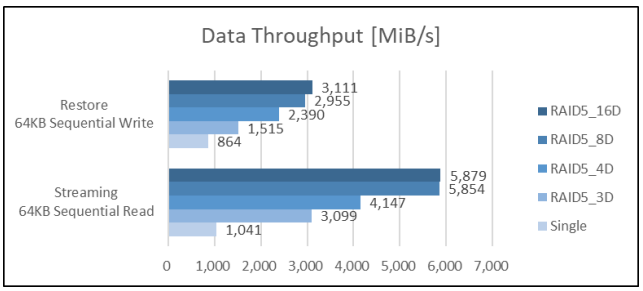


 

# PRAID EP420i
# PCIe option

Up to 8 SAS/SATA drives can be connected directly. The number of drives that can be connected depends on the server configuration. Buffering by cache memory can be expected to improve the performance of the write processing.

In this measurement, in order to see the basic characteristics of the RAID controller, the measurements were performed without using the buffering function by the cache memory.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| HDD: 256 KiB<br>SSD: 64 KiB | Read | No Read Ahead | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #2 | 32 GiB | RAW | Sequential test volume |
| | 64 GiB | RAW | Random test volume |

## *SAS-HDD Connection (Measuring Up to 16 Drives via SAS Switch)*

With RAID0/RAID5, the performance value is proportional to the number of connected drives.

Since the read performance of RAID1 is higher than that of a single drive, the read performance of RAID10 is also improved.

[RAID0]

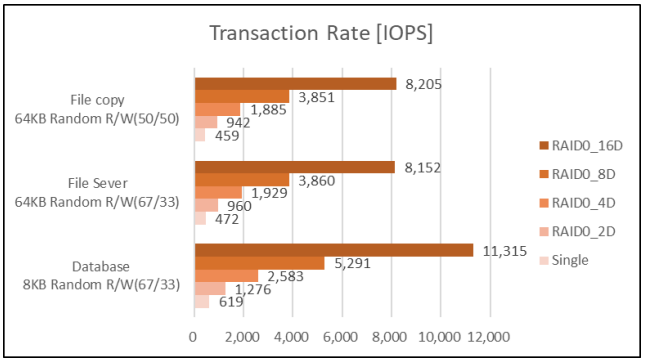


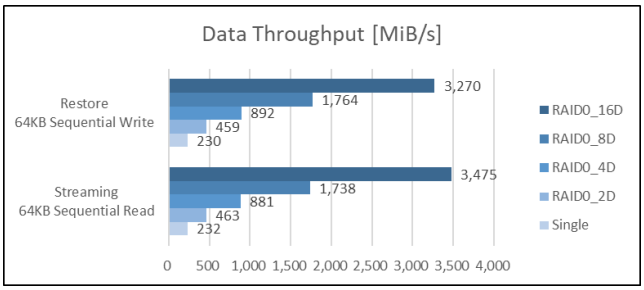


 

[RAID1/10]

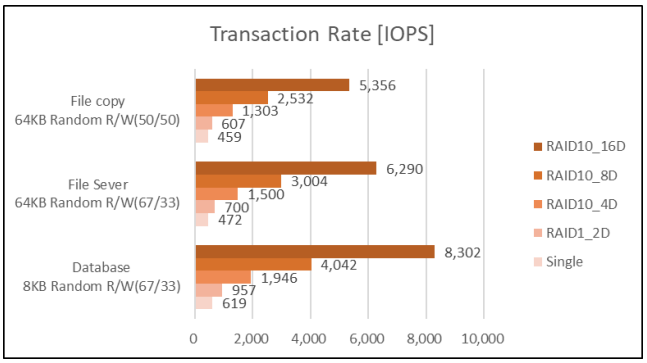


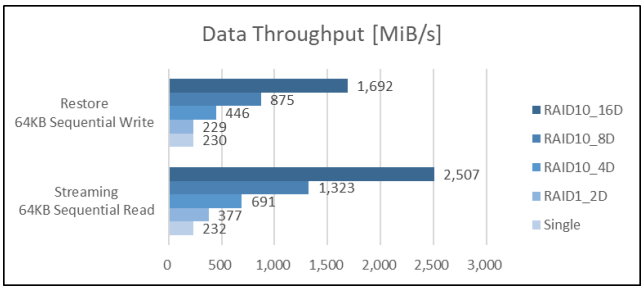


[RAID5]

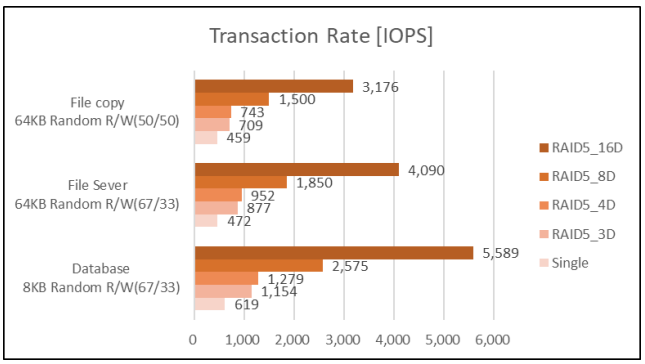


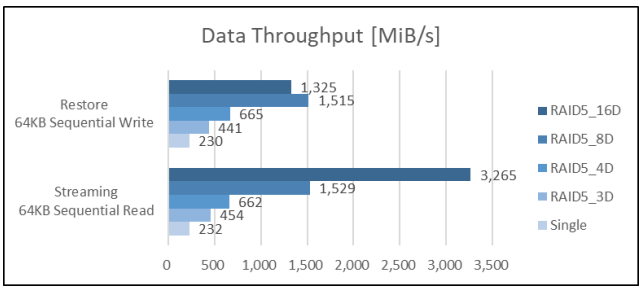


## *SAS-SSD Connection (Measuring Up to 16 Drives via SAS Switch)*

The sequential access performance of 16 drives is almost the same as that of 8 drives. By connecting with 8 drives of about 1,000 MiB/s per drive, the transfer capacity on the drive side can reach about 8,000 MiB/s. However, the theoretical throughput of PCIe Gen3 x8 lane, which is the interface on the CPU side, is 7,512 MiB/s, so it is presumed that the performance limit of the CPU side interface has been reached when eight drives are connected.

In addition, a throughput of about 3,200 MiB/s is obtained with RAID1/10 write performance, and a throughput of about 3,000 MiB/s is obtained with RAID5 write performance.

[RAID0]

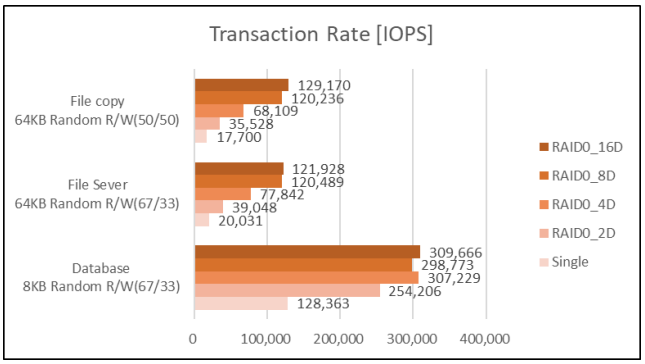


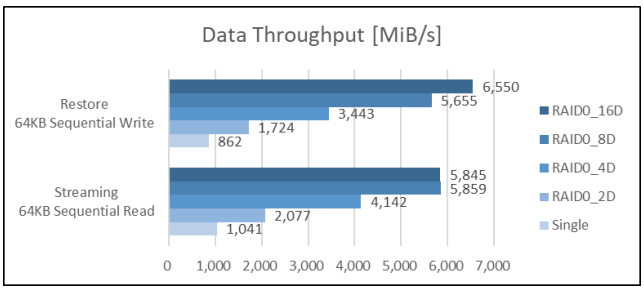


 

## [RAID1/10]

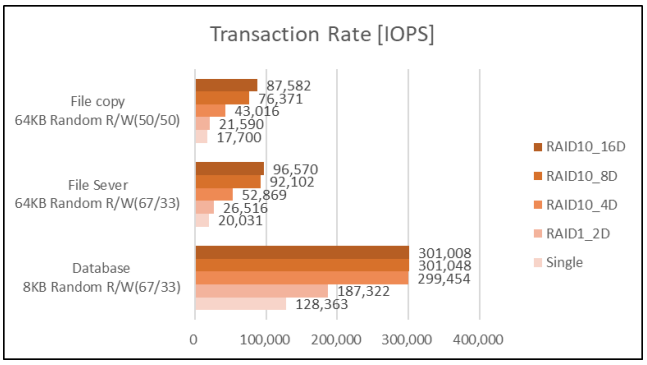


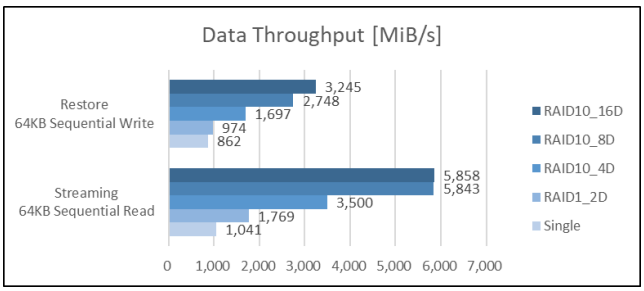


## [RAID5]

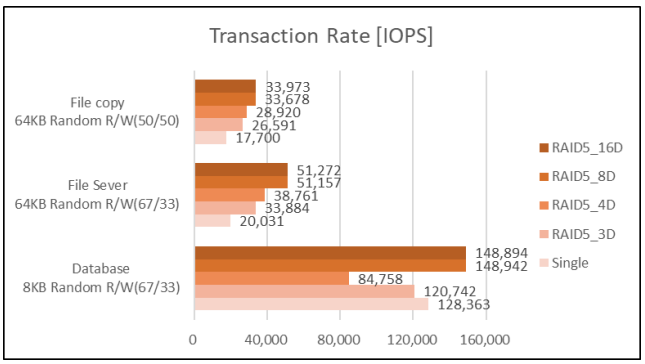


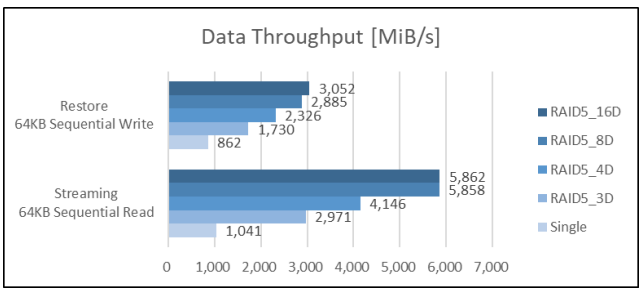


 

# PRAID EP540i
# PCIe option

Up to 16 SAS/SATA drives or 4 NVMe can be connected directly. The number of drives that can be connected depends on the server configuration. Buffering by cache memory can be expected to improve the performance of the write processing.

In this measurement, in order to see the basic characteristics of the RAID controller, the measurements were performed without using the buffering function by the cache memory.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| HDD: 256 KiB<br>SSD/NVMe: 64 KiB | Read | No Read Ahead | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #2 | 32 GiB | RAW | Sequential test volume |
| | 64 GiB | RAW | Random test volume |

## *SAS-HDD Connection*

With RAID0/RAID5, the performance value is proportional to the number of connected drives.

Since the read performance of RAID1 is higher than that of a single drive, the read performance of RAID10 is also improved.

[RAID0]

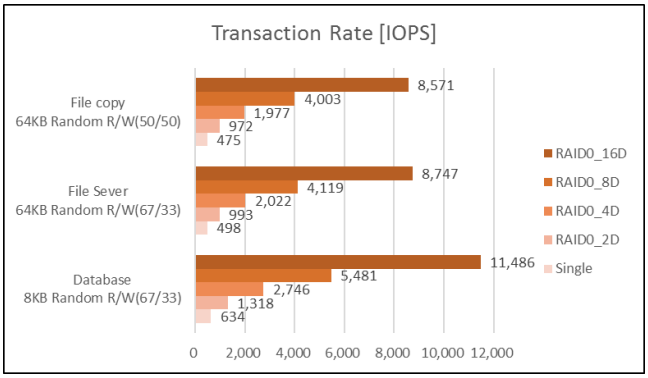


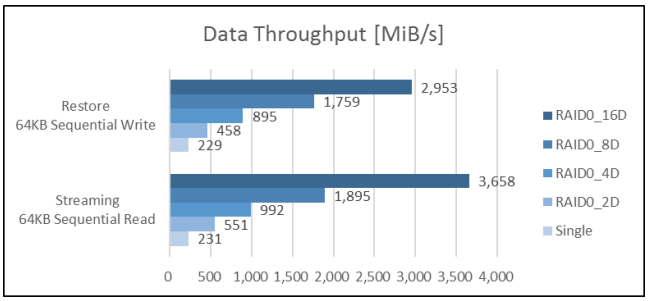


 

[RAID1/10]

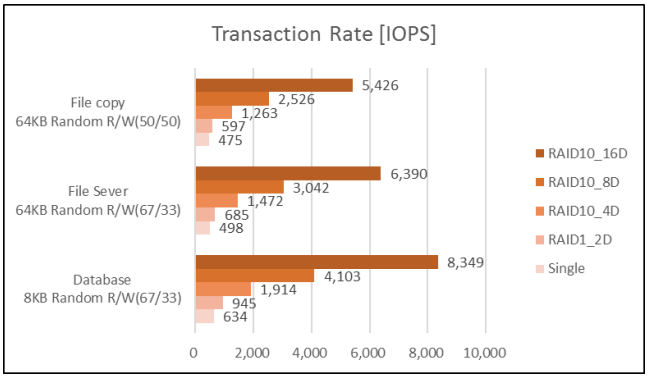


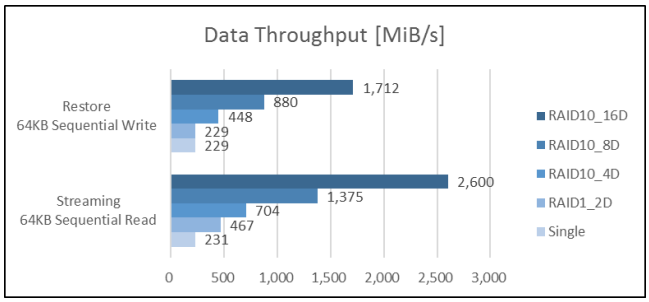


[RAID5]

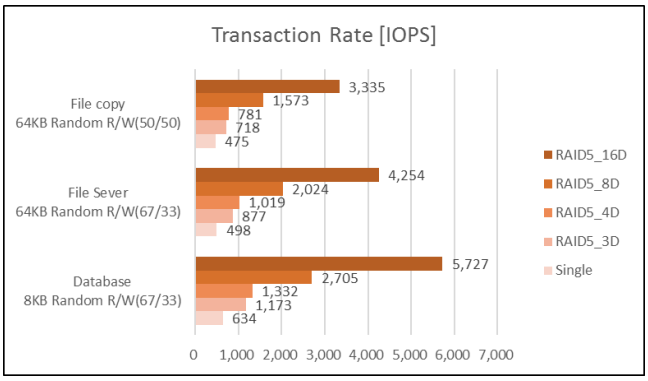


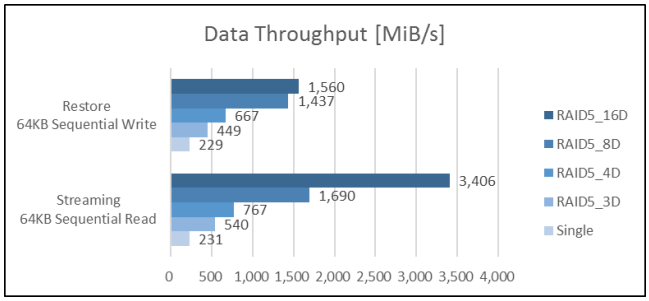


## *SAS-SSD Connection*

The sequential access performance of 16 drives is almost the same as that of 8 drives. By connecting with 8 drives of about 1,000 MiB/s per drive, the transfer capacity on the drive side can reach about 8,000 MiB/s. However, the theoretical throughput of PCIe Gen3 x8 lane, which is the interface on the CPU side, is 7,512 MiB/s, so it is presumed that the performance limit of the CPU side interface has been reached when eight drives are connected.

In addition, a throughput of about 4,100 MiB/s is obtained with RAID1/10 write performance, and a throughput of about 3,400 MiB/s is obtained with RAID5 write performance.

[RAID0]

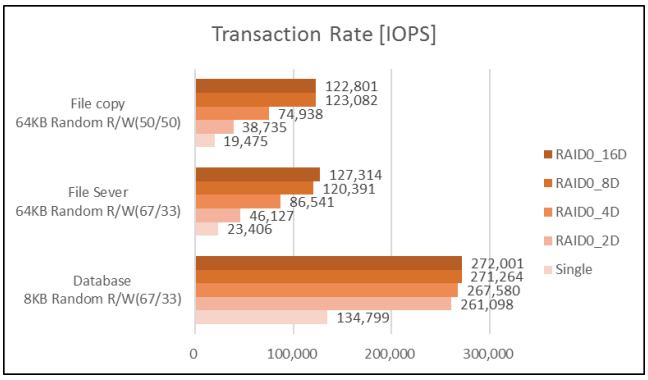


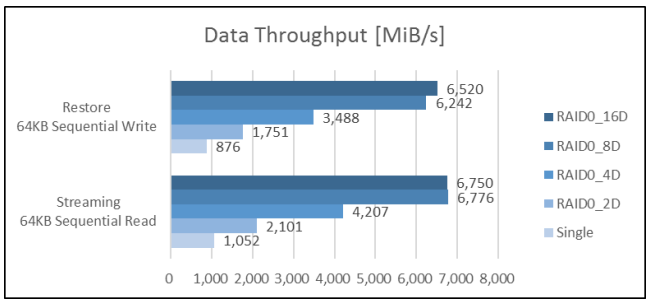


 

[RAID1/10]

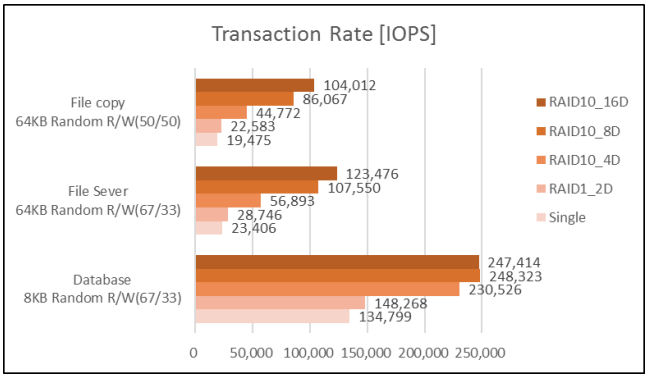


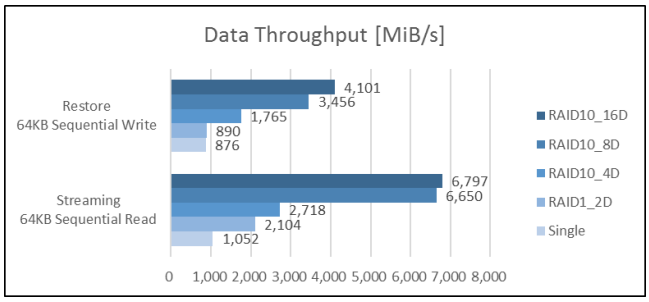


[RAID5]

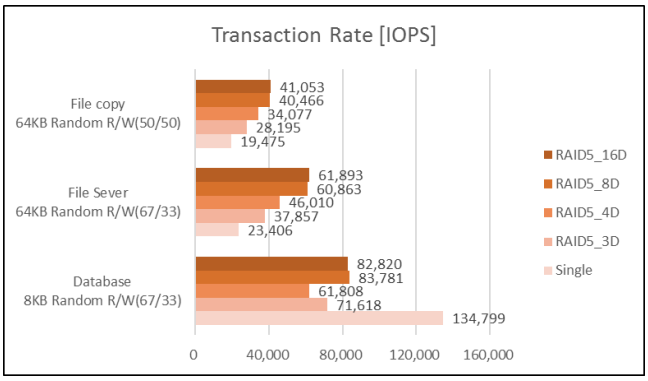


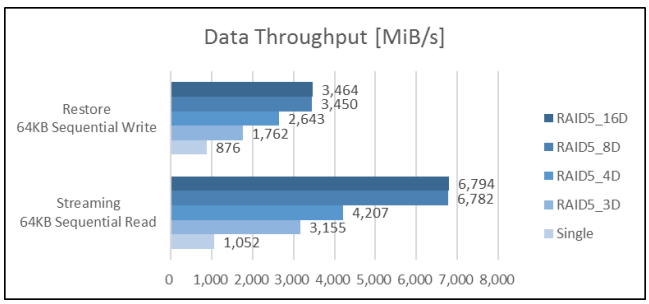


## *NVMe Connection*

Since the read performance almost reached the limit when 2 drives are connected, the performance value when 4 drives were connected was saturated.

Since the performance value of the database pattern, which is a random access, was measured with "# of Worker Threads = 1", the latency of the transfer process becomes a bottleneck, and the transaction rate has peaked.

[RAID0]

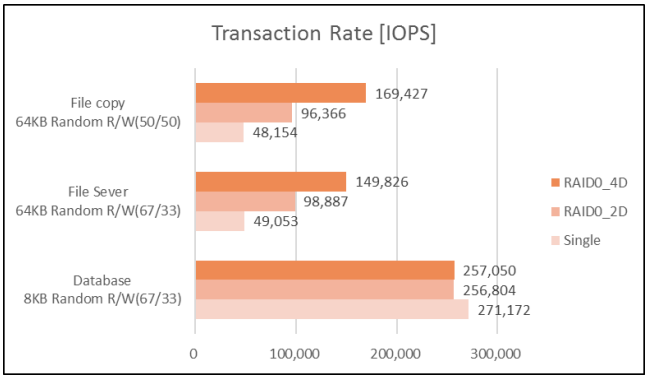


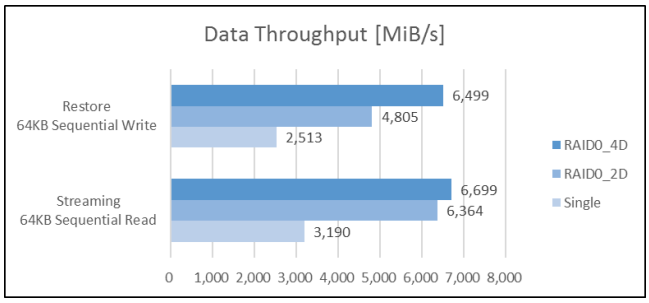


 

## [RAID1/10]

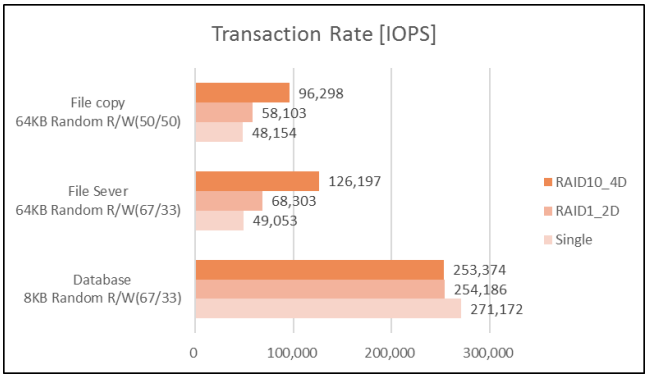


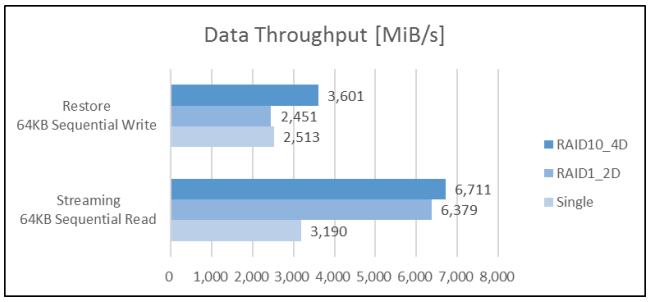


## [RAID5]

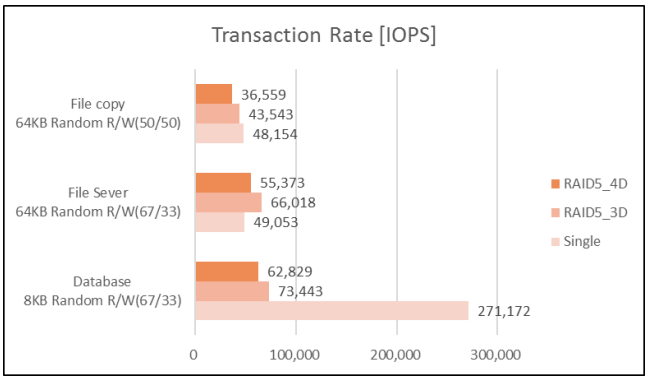


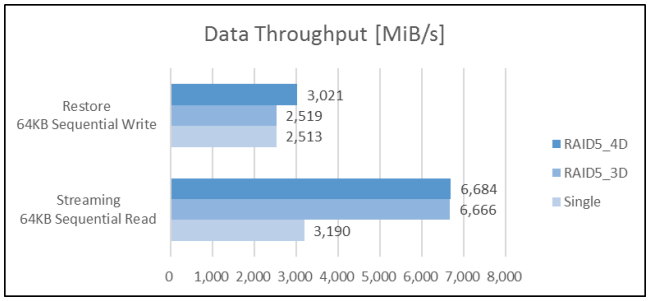


 

# PRAID EP640i
# PCIe option

The host bus supports PCIe 4.0, and the data bus bandwidth is improved. Up to 8 SAS/SATA drives can be directly connected to the RAID controller. The number of drives that can be connected depends on the server configuration. Buffering by cache memory can be expected to improve the performance of the write processing.

In this measurement, in order to see the basic characteristics of the RAID controller, the measurement is performed without using the buffering function by the cache memory.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| HDD: 256 KiB<br>SSD: 64 KiB | Read | No Read Ahead | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #2 | 32 GiB | RAW | Sequential test volume |
| | 64 GiB | RAW | Random test volume |

## *SAS-HDD Connection (Measuring Up to 16 Drives via SAS Switch)*

With RAID0/RAID5, the performance value is proportional to the number of connected drives.

The sequential performance of RAID1 is almost the same as the performance of a single drive. RAID10 has a performance proportional to the performance of RAID1 (double for four drives and 4 times for eight drives).

[RAID0]

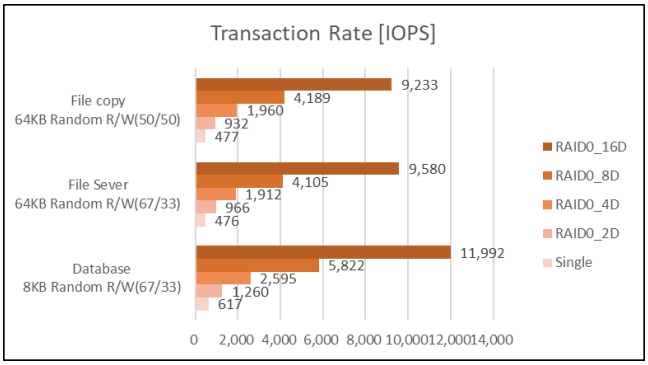


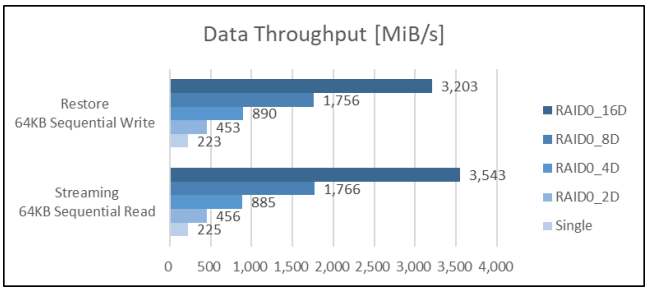


 

[RAID1/10]

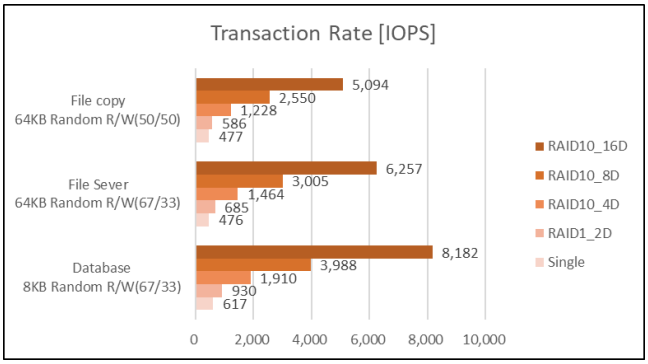


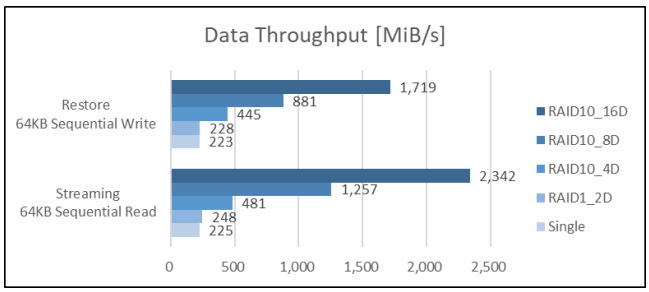


[RAID5]

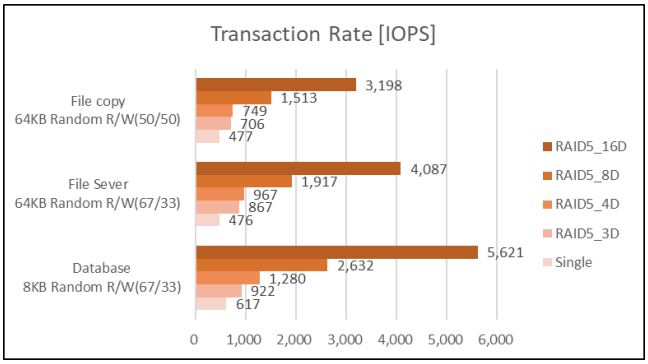


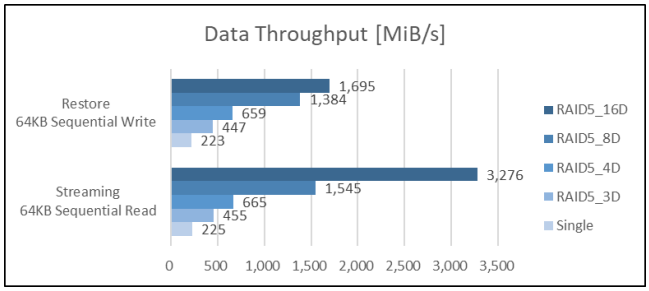


## *SAS-SSD Connection (Measuring Up to 16 Drives via SAS Switch)*

The sequential access performance of 16 drives is almost the same as that of 8 drives. Since the number of ports of drive-side is 8 ports, it is assumed that the drive-side throughput limit of about 8,000 MiB/s was reached at the time of 8 drives were connected. Since the performance limit has been reached, increasing the number of drives will not improve performance.

In addition, a throughput of about 4,000 MiB/s is obtained with RAID1/10 write performance, and a throughput of about 5,200 MiB/s is obtained with RAID5 write performance.

[RAID0]

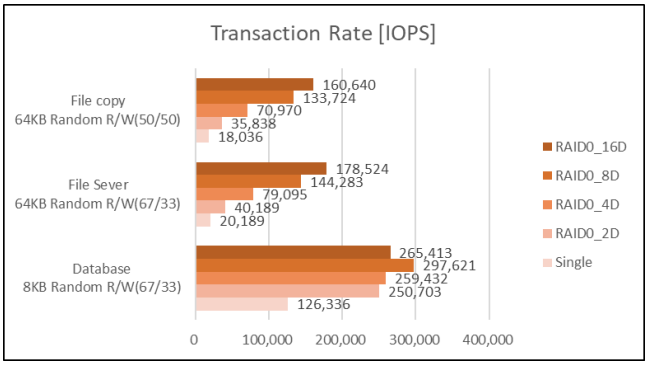


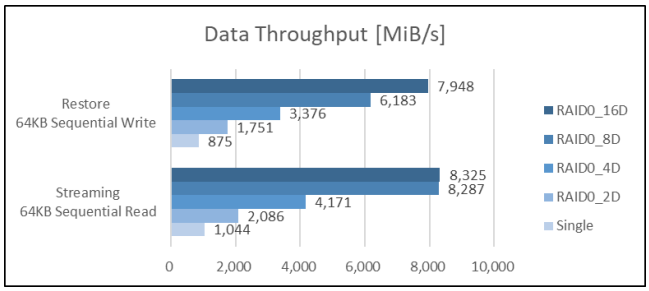


 

## [RAID1/10]

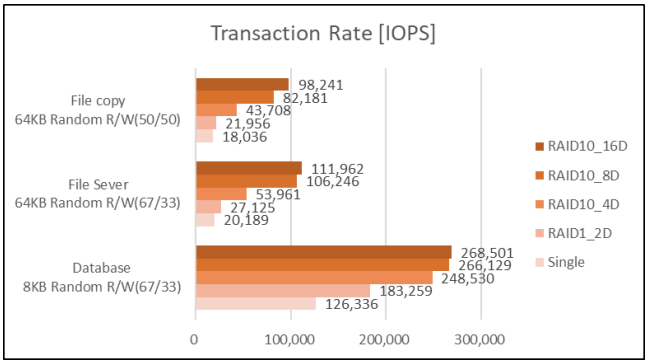


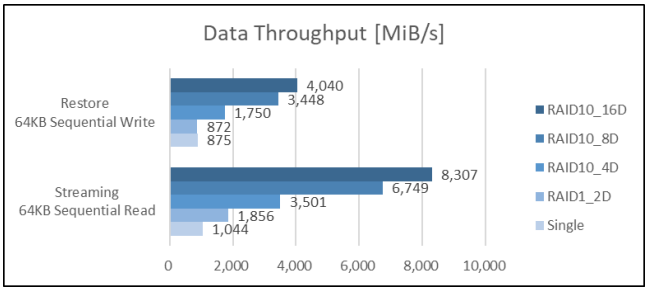


## [RAID5]

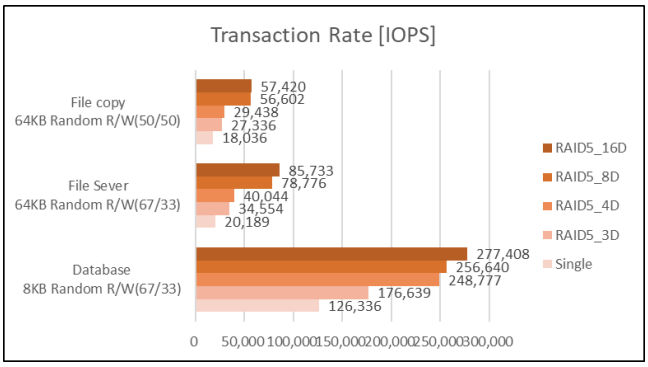


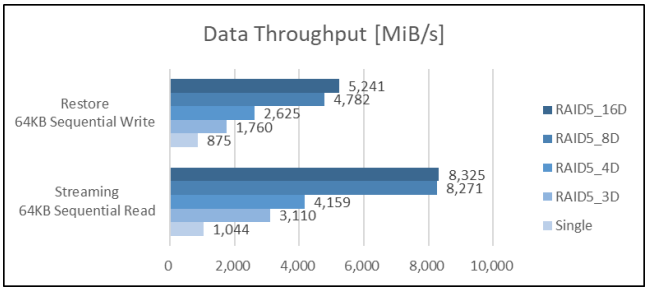


 

# PRAID EP680i
# PCIe option

The host bus supports PCIe 4.0, and the data bus bandwidth is improved. Up to 16 SAS/SATA drives or 4 NVMe drives can be directly connected to the RAID controller. The number of drives that can be connected depends on the server configuration. Buffering by cache memory can be expected to improve the performance of the write processing.

In this measurement, in order to see the basic characteristics of the RAID controller, the measurement is performed without using the buffering function by the cache memory.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| HDD: 256 KiB<br>SSD/NVMe: 64 KiB | Read | No Read Ahead | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #2 | 32 GiB | RAW | Sequential test volume |
| | 64 GiB | RAW | Random test volume |

## *SAS-HDD Connection*

With RAID0/RAID5, the performance value is proportional to the number of connected drives.

The sequential performance of RAID1 is almost the same as the performance of a single drive. RAID10 has a performance proportional to the performance of RAID1 (double for four drives and 4 times for eight drives).

[RAID0]

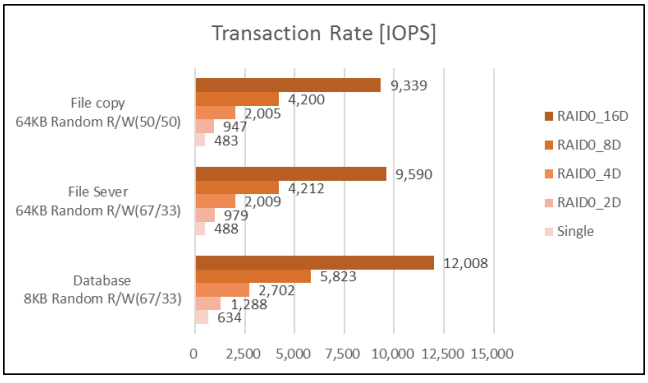


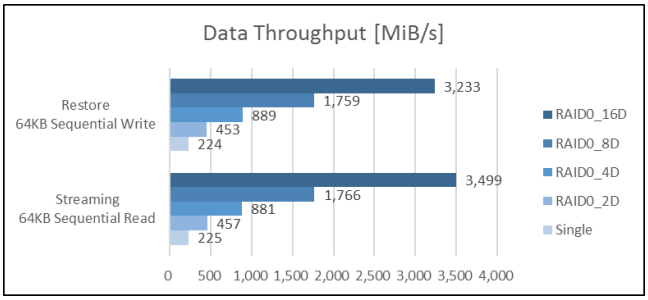


 

[RAID1/10]

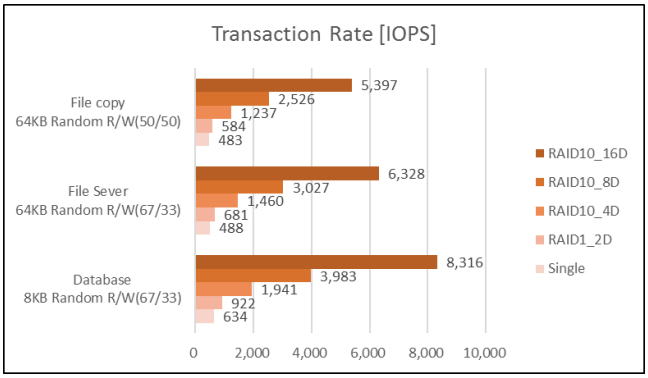


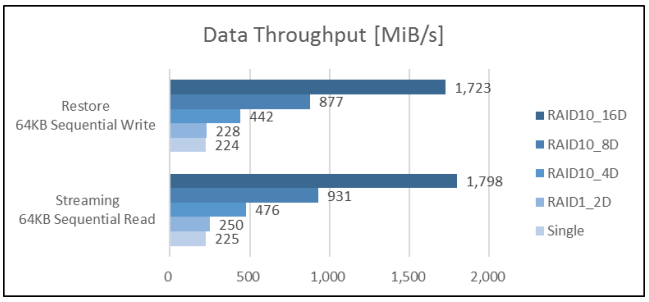


[RAID5]

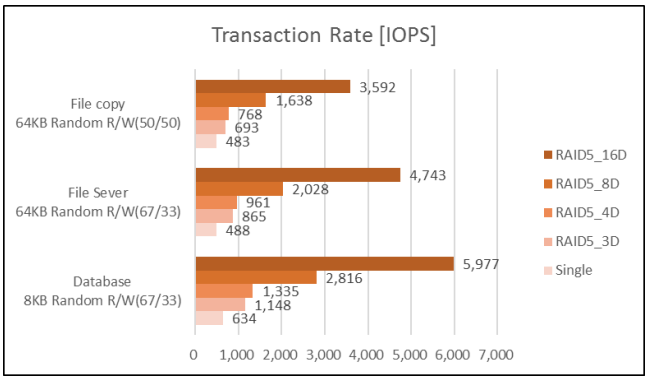


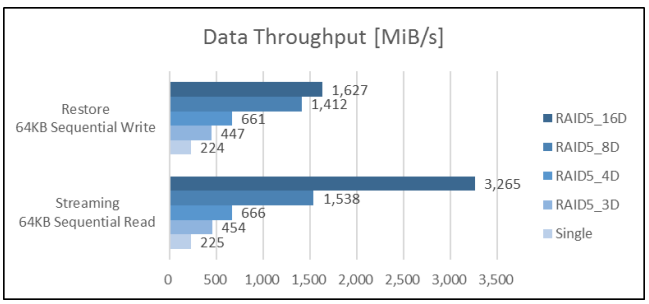


## *SAS-SSD Connection*

Since the interface on the CPU side is PCIe Gen4, the bandwidth is improved and the performance value according to the number of connected devices is obtained. In addition, write performance of RAID1/10/5 shows that the processing power of the controller is improved, with a throughput of about 6,500 MiB/s for RAID1/10 and a throughput of about 4,700 MiB/s for RAID5.

[RAID0]

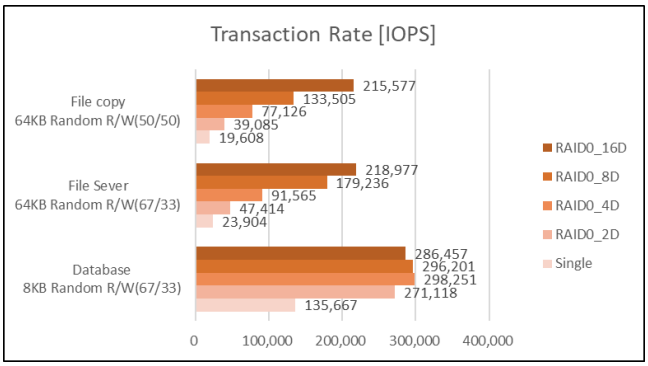


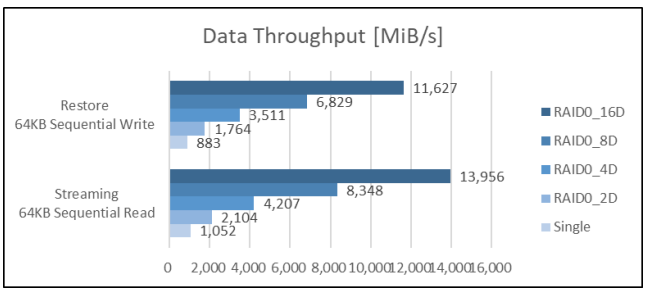


 

[RAID1/10]

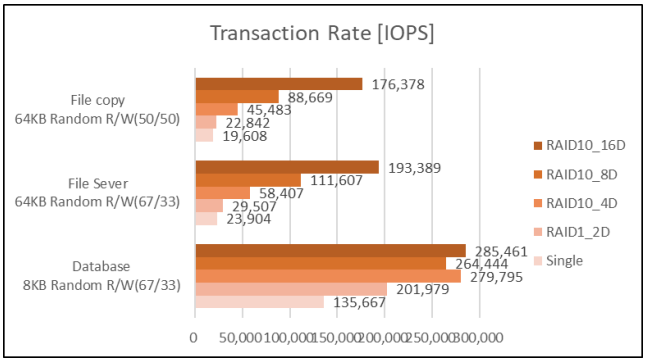


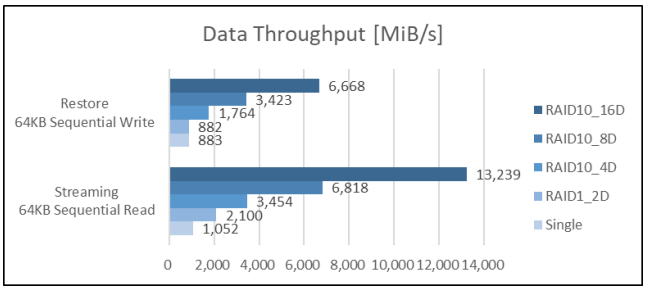


[RAID5]

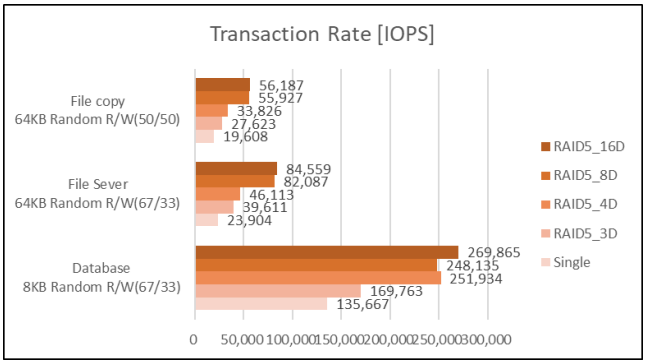


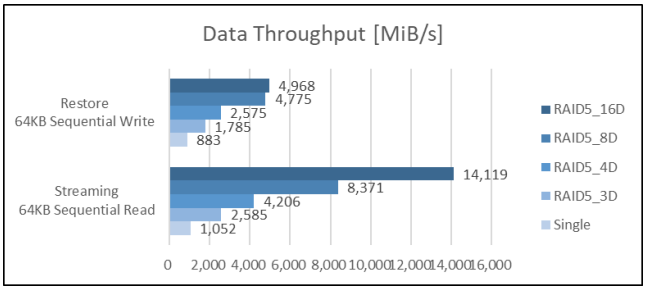


## *NVMe Connection*

Since the CPU side interface is PCIe Gen4, the bandwidth is improved, and the performance value increased according to the number of connected drives. However, when RAID0 of 4 drives are connected, the read performance reaches the performance limit of the controller and remains at about 10,000 MiB/s. With RAID1, the read is twice that of a single drive, the write is the same, and the expected performance is obtained. Write performance of RAID5 shows that the processing power of the controller is improved, and the throughput of about 4,100 MiB/s is obtained. (Approximately 3,000 MiB/s for PRAID EP540i)

Since the performance value of the database pattern, which is a random access, is measured by one worker, the latency of the transfer process becomes a bottleneck, and the transaction rate has peaked.

[RAID0]

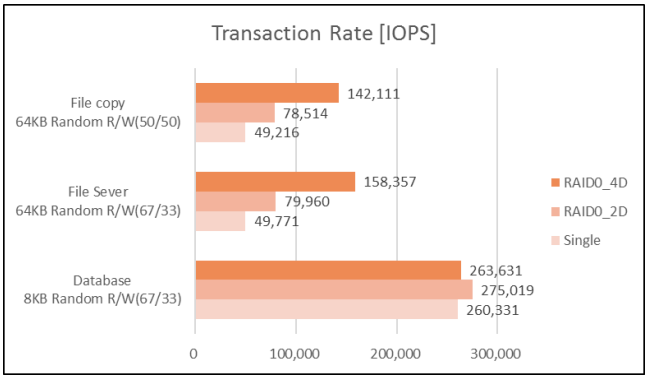


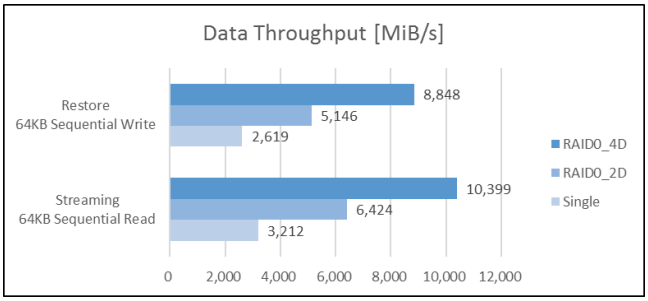


 

## [RAID1/10]

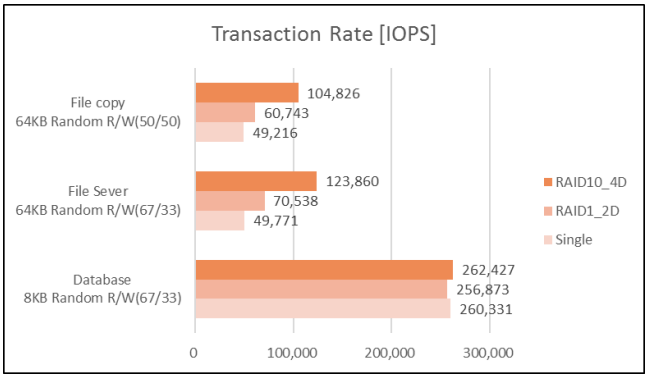


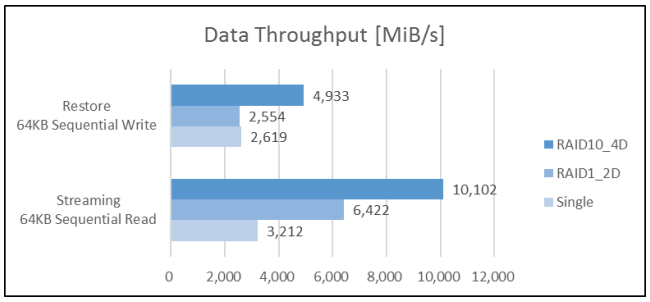


## [RAID5]

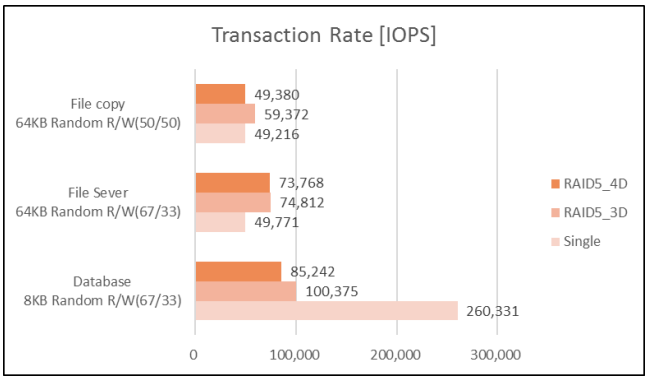


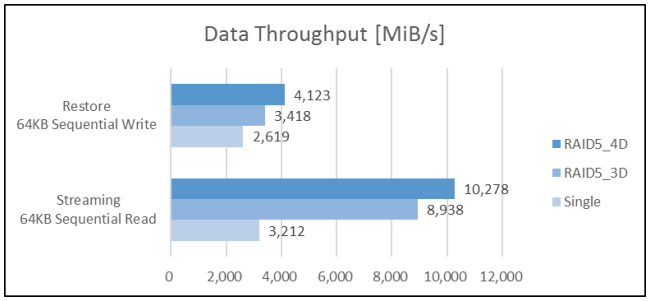


 

# PRAID EP3254-8i
# PCIe option

The host bus supports PCIe 4.0, and up to 8 SAS/SATA drives can be directly connected to the RAID controller. The number of drives that can be connected depends on the server configuration. Buffering by cache memory can be expected to improve the performance of the write processing.

In this measurement, in order to see the basic characteristics of the RAID controller, the measurement is performed without using the buffering function by the cache memory.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| HDD: 256 KiB<br>SSD: 64 KiB | Read | - | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #1 | 64 GiB | RAW | Sequential test volume<br>Random test volume |

## *SAS-HDD Connection (Measuring Up to 16 Drives via SAS Switch)*

With RAID0, the performance value is proportional to the number of connected drives.

When HDD is connected, sequential write of RAID1/10/5 becomes a low performance value, so it is recommended to work with write-back enabled by using a flash backup unit.

[RAID0]

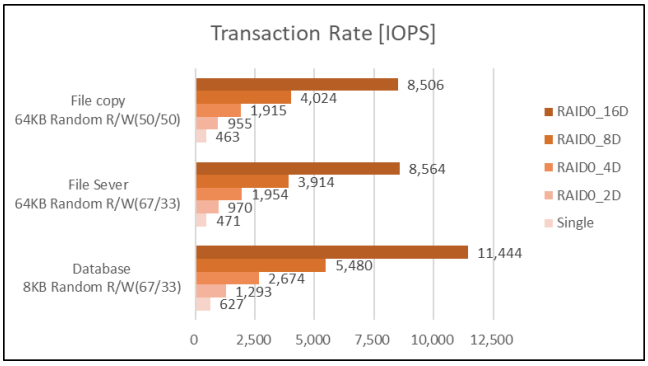


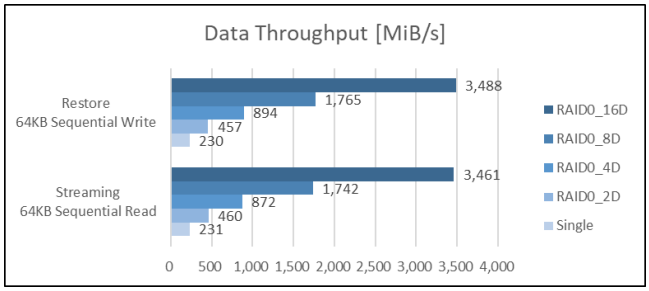


 

## [RAID1/10]

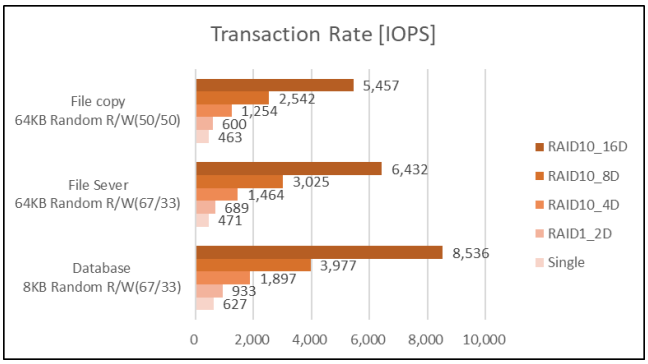


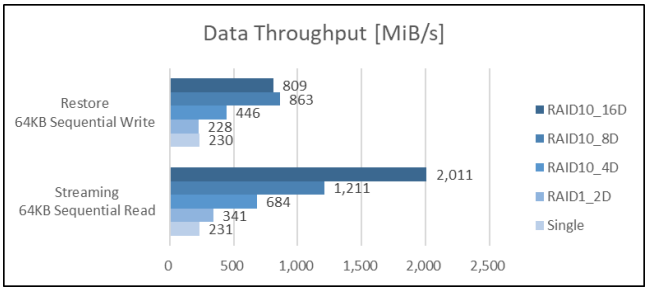


## [RAID5]

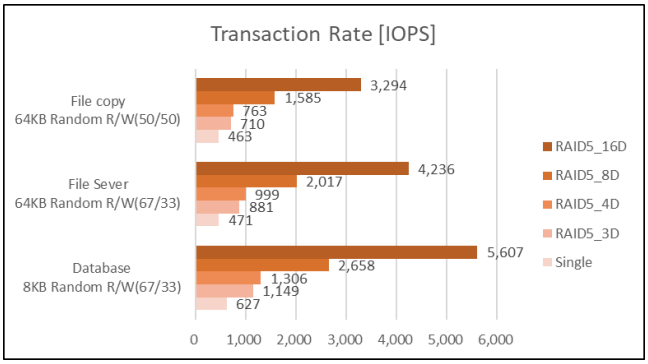


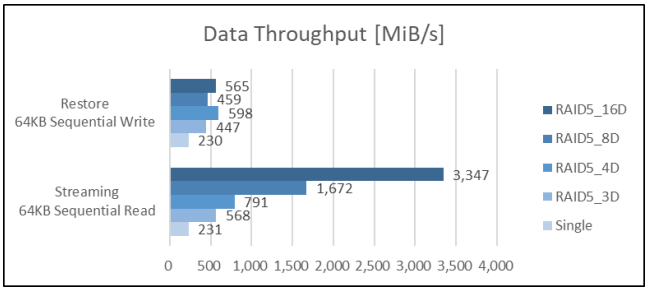


## Reference: Performance when setting write-back

## [RAID1/10]

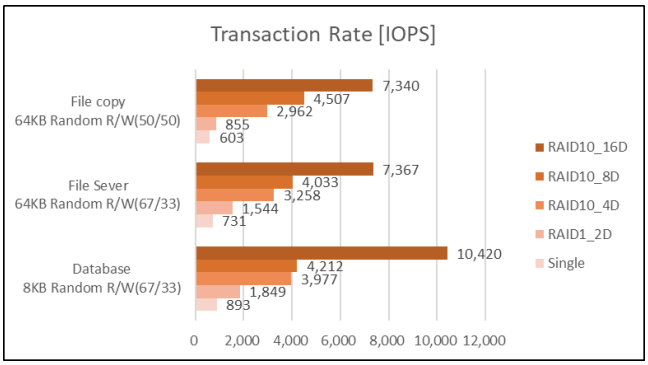


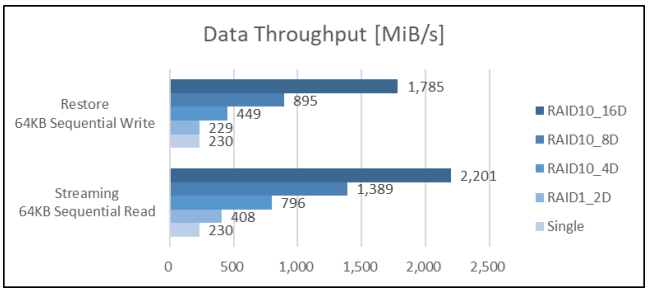


## [RAID5]

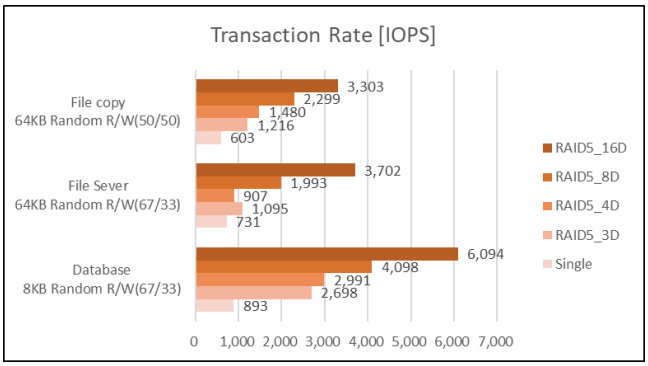


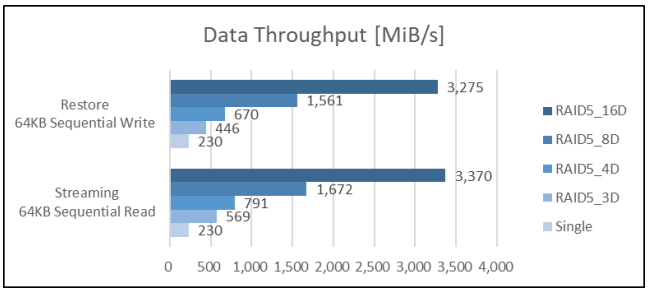


 

## *SAS-SSD Connection (Measuring Up to 16 Drives via SAS Switch)*

The sequential access performance of 16 drives is almost the same as that of 8 drives. Since the number of ports of drive-side is 8 ports, it is assumed that the drive-side throughput limit of about 8,000 MiB/s was reached at the time of 8 drives were connected. Since the performance limit has been reached, increasing the number of drives will not improve performance.

In addition, a throughput of about 3,600 MiB/s is obtained with RAID1/10 write performance, and a throughput of about 5,000 MiB/s is obtained with RAID5 write performance.

[RAID0]

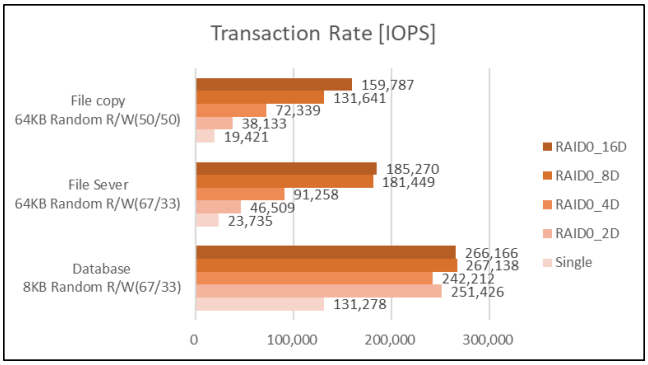


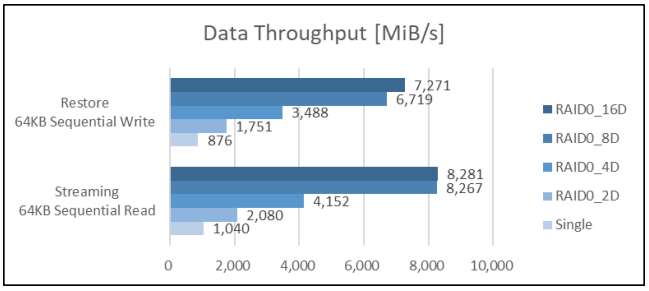


[RAID1/10]

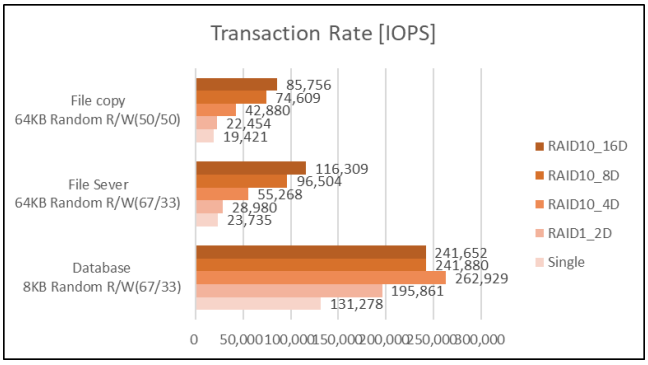


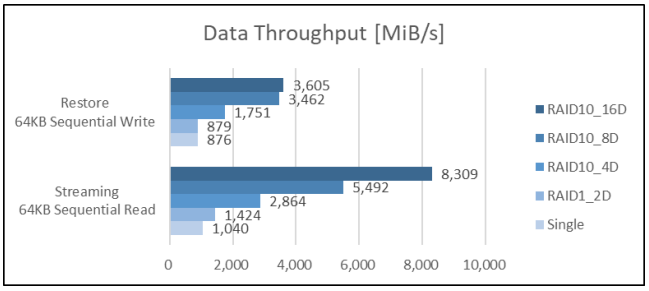


[RAID5]

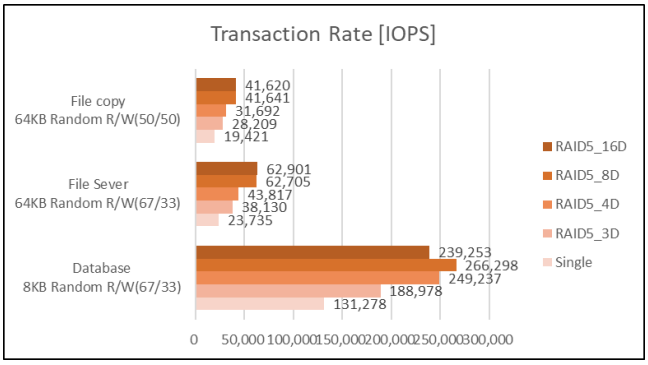


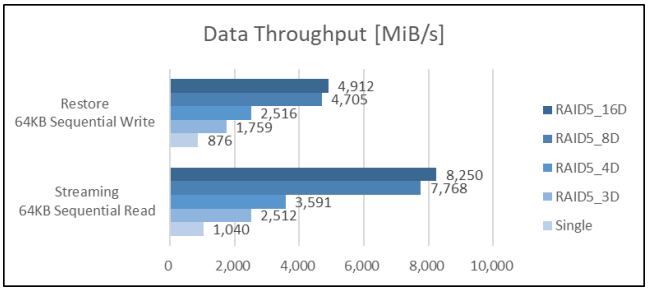


 

# PRAID EP3258-16i
# PCIe option

The host bus supports PCIe 4.0, and up to 16 SAS/SATA drives or 4 NVMe drives can be directly connected to the RAID controller. The number of drives that can be connected depends on the server configuration. Buffering by cache memory can be expected to improve the performance of the write processing.

In this measurement, in order to see the basic characteristics of the RAID controller, the measurement is performed without using the buffering function by the cache memory.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|
| HDD: 256 KiB<br>SSD: 64 KiB | Read | - | Disabled |
| | Write | Write Through | |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #1 | 64 GiB | RAW | Sequential test volume<br>Random test volume |

## *SAS-HDD Connection*

With RAID0, the performance value is proportional to the number of connected drives.

When HDD is connected, sequential write of RAID1/10/5 becomes a low performance value, so it is recommended to work with write-back enabled by using a flash backup unit.

[RAID0]

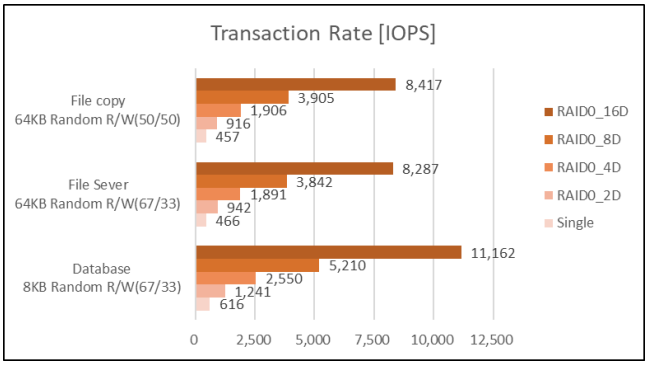


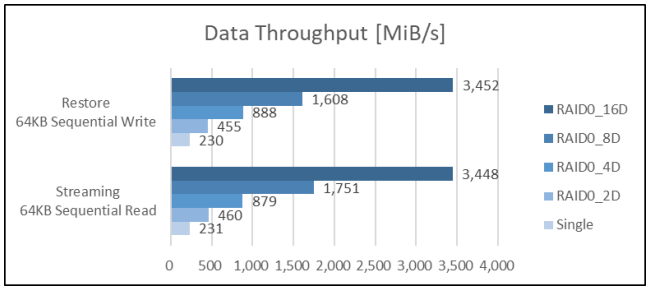


 

## [RAID1/10]

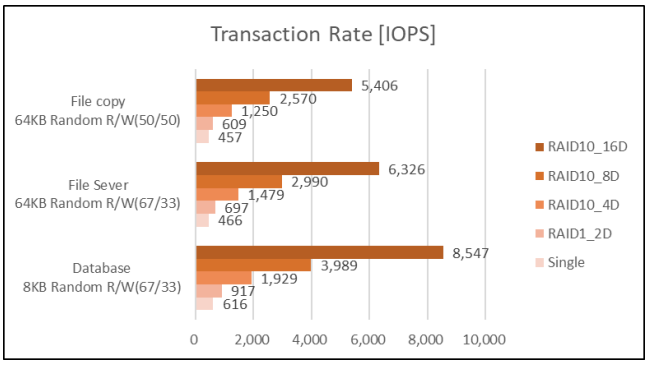


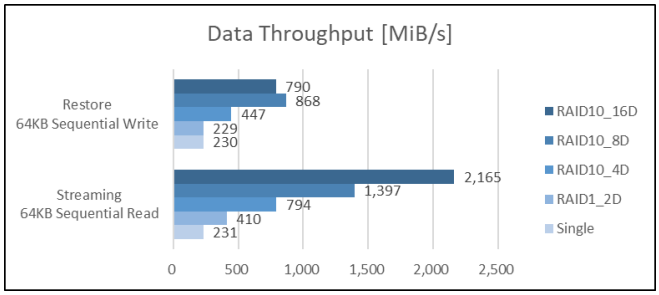


## [RAID5]

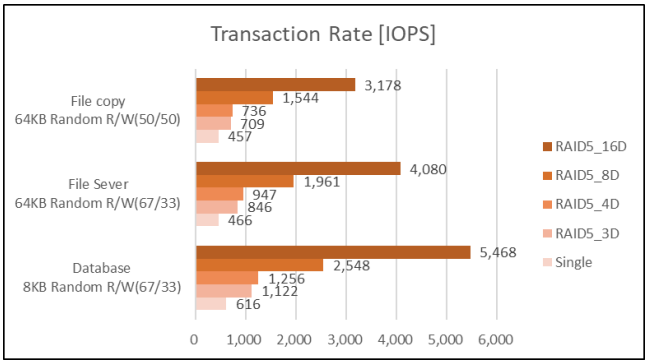


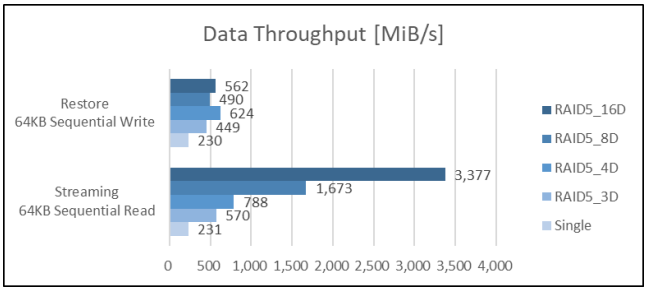


## Reference: Performance when setting write-back

## [RAID1/10]

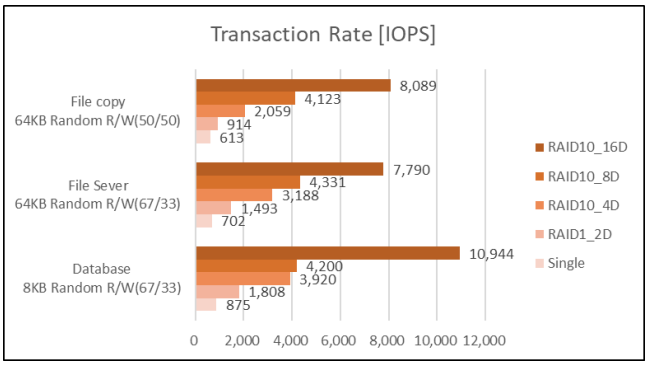


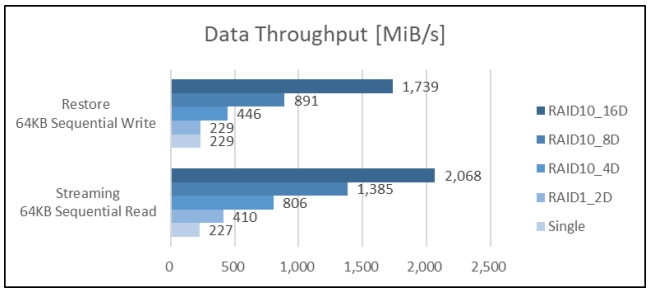


## [RAID5]

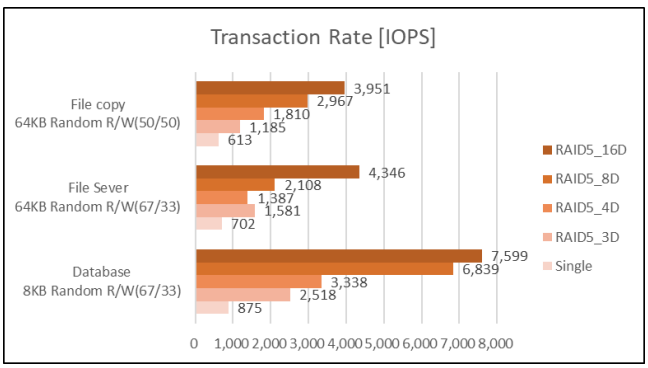


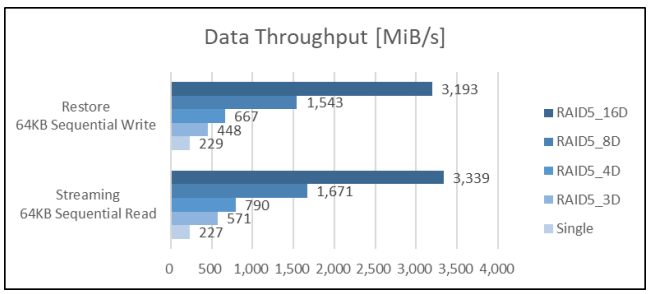


 

## *SAS-SSD Connection*

Since the interface on the CPU side is PCIe Gen4, the bandwidth is improved and the performance value according to the number of connected devices is obtained. In addition, write performance of RAID1/10/5 shows that the processing power of the controller is improved, with a throughput of about 6,300 MiB/s for RAID1/10 and a throughput of about 5,000 MiB/s for RAID5.

[RAID0]

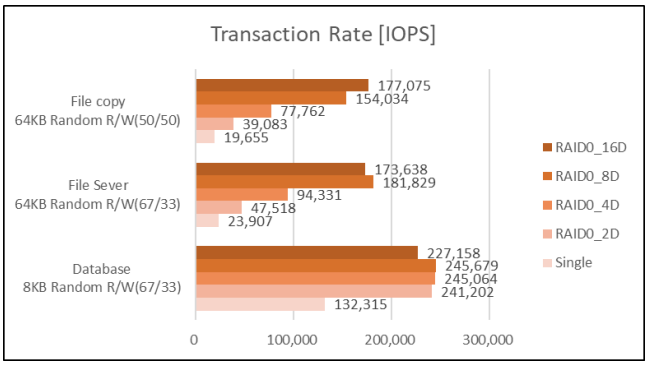


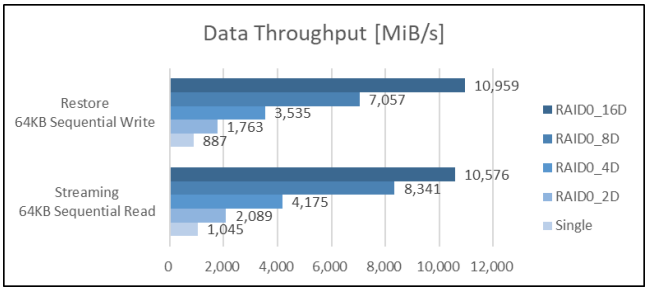


[RAID1/10]

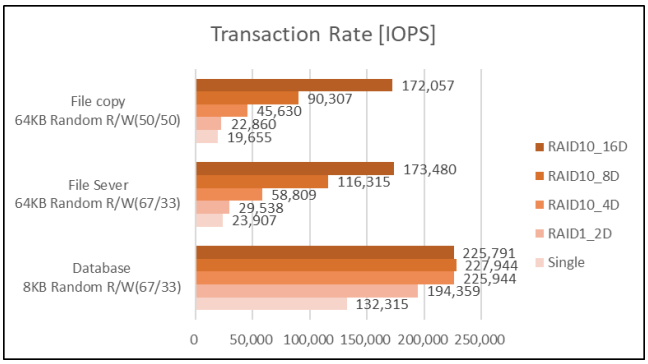


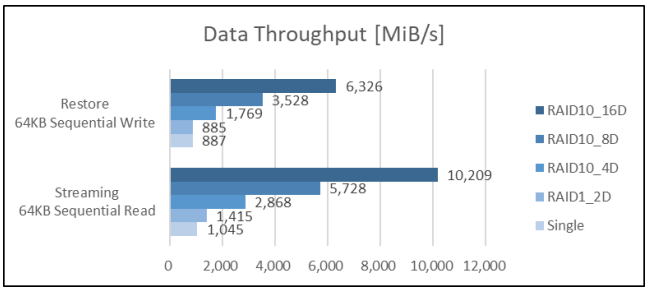


[RAID5]

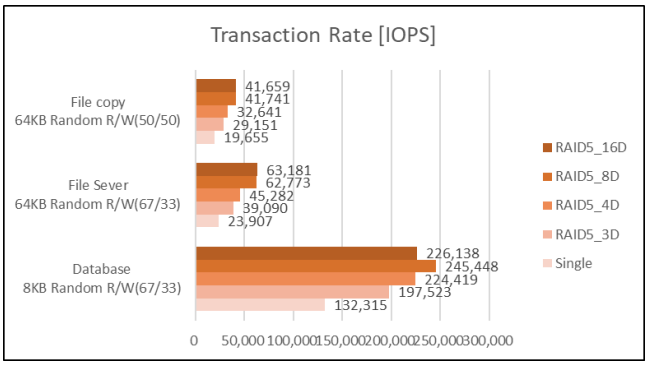


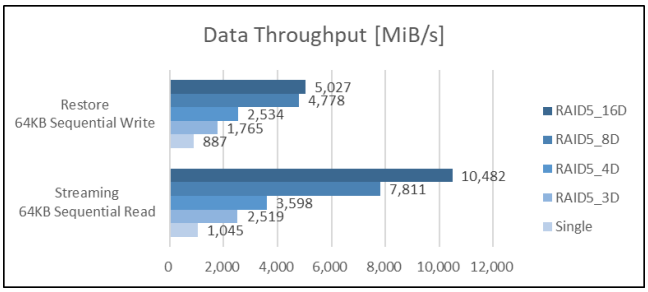


 

## *NVMe Connection*

Since the CPU side interface is PCIe Gen4, the bandwidth is improved, and the performance value increased according to the number of connected drives. With RAID0/1/10, the performance value is proportional to the number of connected drives. The write performance of RAID5 shows that the processing power of the controller, with a throughput of about 4,200 MiB/s.

Since the performance value of the database pattern, which is a random access, is measured by one worker, the latency of the transfer process becomes a bottleneck, and the transaction rate has peaked.

[RAID0]

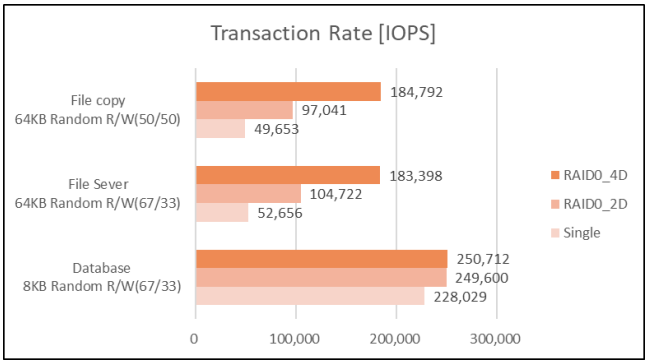


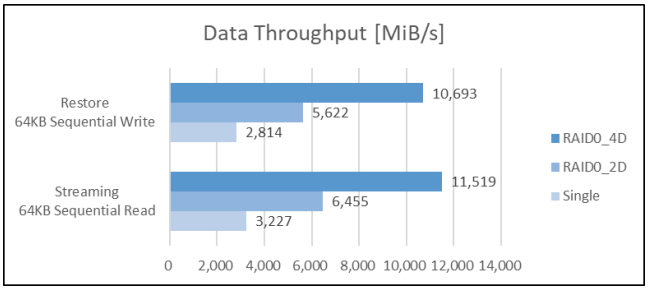


[RAID1/10]

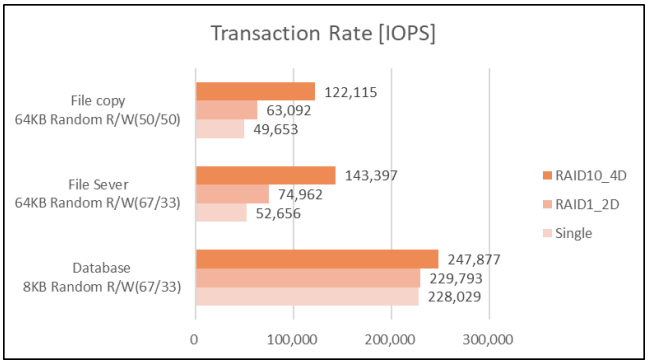


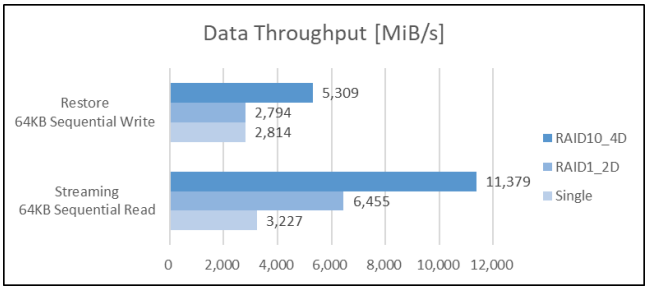


[RAID5]

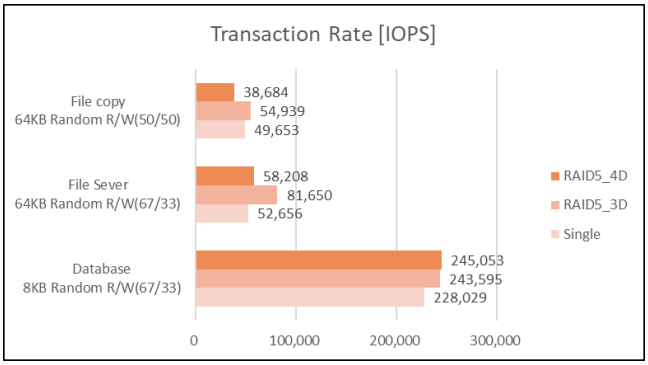


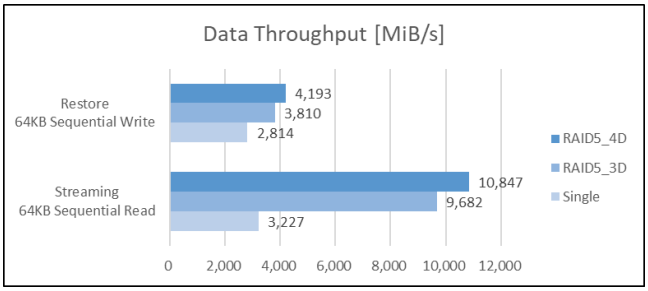


 

# Effect of the controller cache

Using the controller cache can improve random access performance.

WT(Write Through): This mode waits for a completion response from the drive during write access. This mode does not use the controller cache.

WB(Write Back): This mode uses the controller cache. Write access is cached in the controller cache and immediately returns a completion response to the application. The drive will be written to later.

Instead of waiting for write access to the drive to complete, applications can perform the following processing. This can improve throughput. If a flash backup unit is not connected and the controller cache data cannot be protected, it automatically switches to WT behavior.

AWB(Always Write Back): This mode forces the controller cache to be used. If no flash backup unit is connected, the controller cache data is not protected.

The following RAID settings were used to measure the access performance of each RAID controller.

Detailed data can be found in Appendix 2: Cache effect measurement results.

## *RAID Settings and Test Volume*

The measurements were performed by the following settings and logical volumes.

| Controller | Drive | Cache Setting | Strip Size | Controller Cache policy | | Disk Cache |
|---|---|---|---|---|---|---|
| | | | | Read | Write | |
| PSAS CP2100-8i | HDD<br>SSD | - | 256KiB | - | - | Disable |
| PSAS CP2200-16i | HDD | - | 256KiB | - | - | Disable |
| | SSD | | 64KiB | | | |
| PRAID CP400i<br>PRAID CP500i<br>PRAID CP600i | HDD<br>SSD | - | 64KiB | - | - | Disable |
| PRAID EP400i<br>PRAID EP420i<br>PRAID EP520i<br>PRAID EP540i<br>PRAID EP640i<br>PRAID EP680i | HDD | WT | 256KiB | Read Ahead | Write Through | Disable |
| | | AWB | | | Always Write Back | Enable |
| | SSD | WT | 64KiB | No Read Ahead | Write Through | Disable |
| | | AWB | | | Always Write Back | Enable |
| PRAID EP3254-8i<br>PRAID EP3258-16i | HDD | WT | 256KiB | - | Write Through | Disable |
| | | AWB | | | Always Write Back | Enable |
| | SSD | WT | 64KiB | - | Write Through | Disable |
| | | AWB | | | Always Write Back | Enable |

| Test Volume | Volume Size | File system | Note |
|---|---|---|---|
| #1 | 64 GiB | RAW | Sequential test volume<br>Random test volume |

 

## *SAS-HDD Connection*

The figure below shows comparison data for a single drive and four drives at each RAID level.

Enabling cache improves throughput. RAID 5 is particularly effective and has significantly improved throughput.

[Single drive]

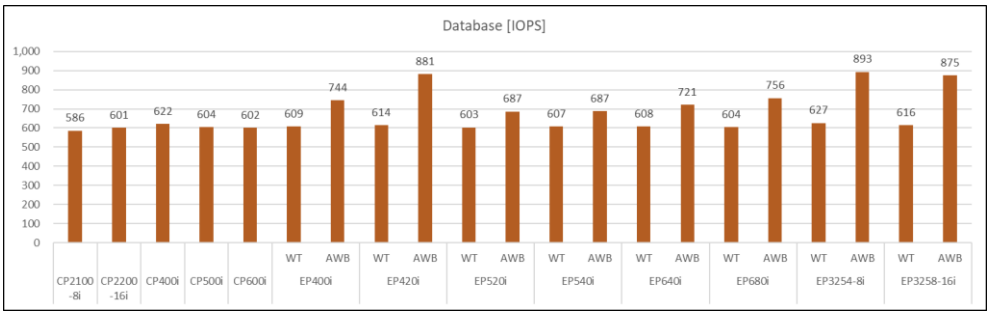


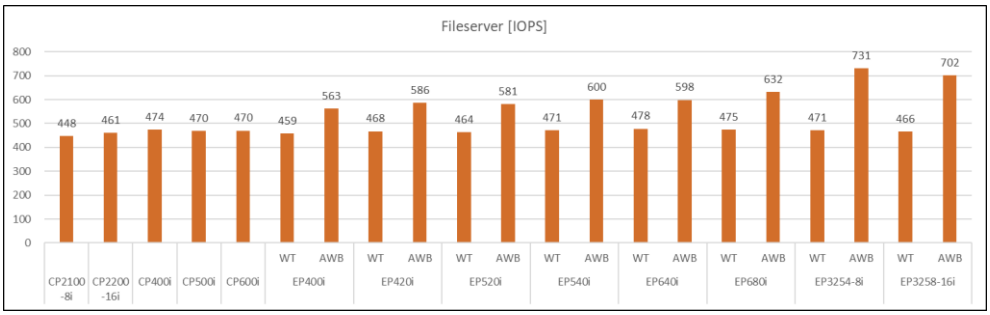


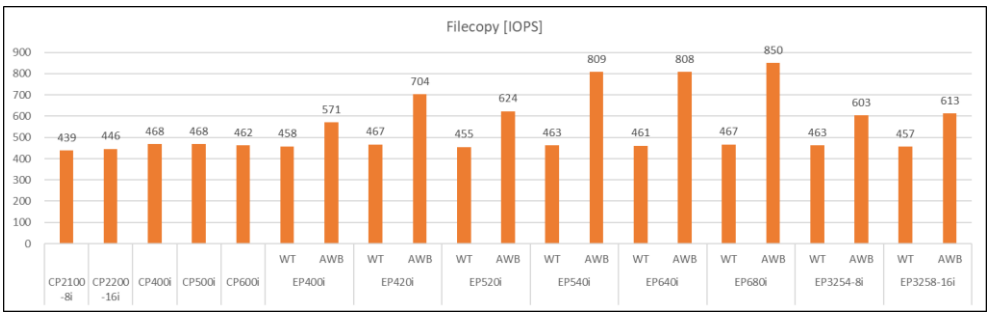


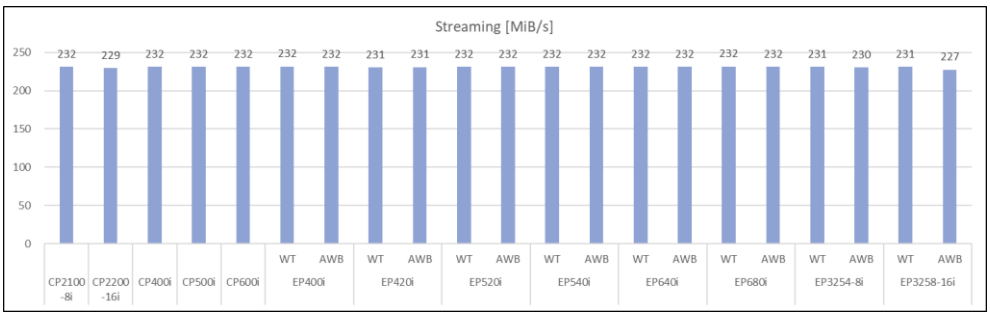



 

## [RAID0 4x drives]

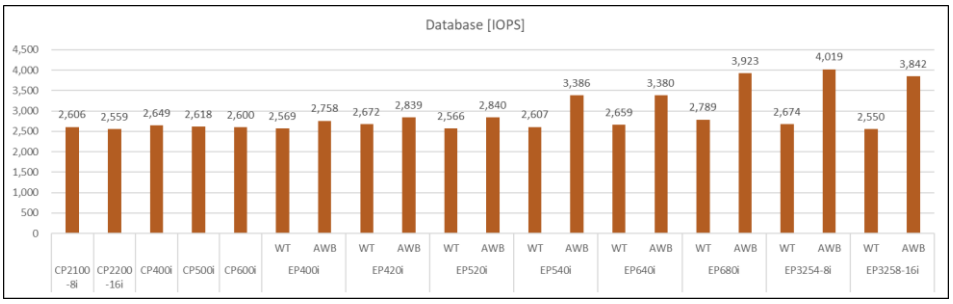


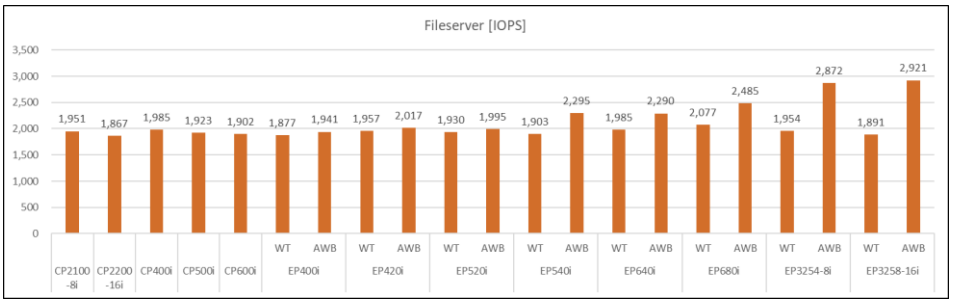


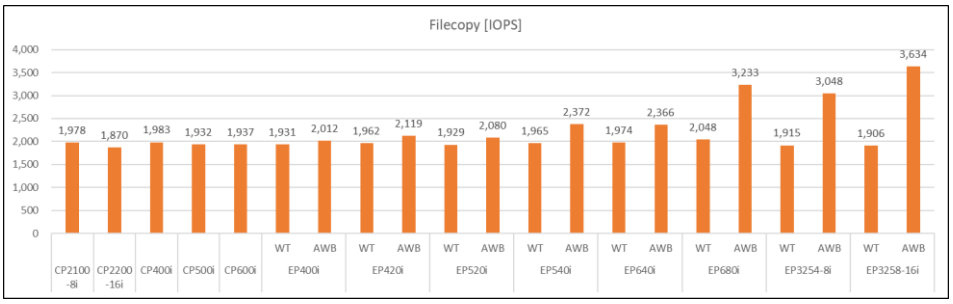


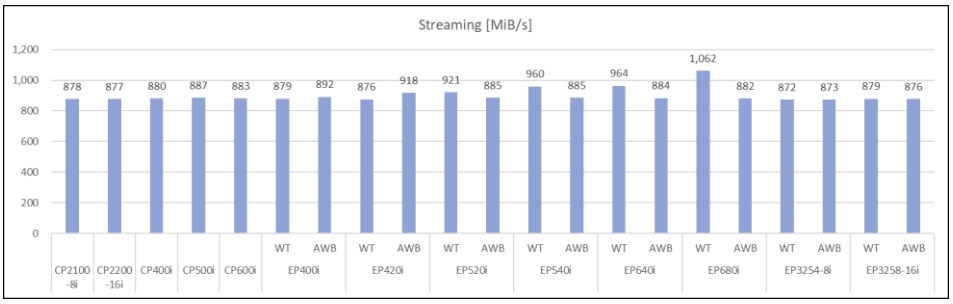


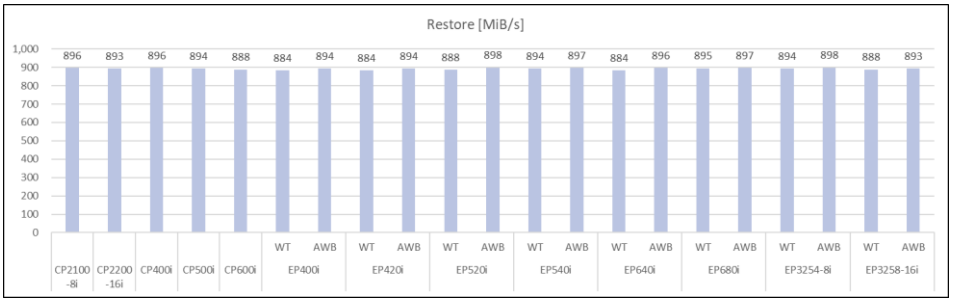


 

## [RAID10 4x drives]

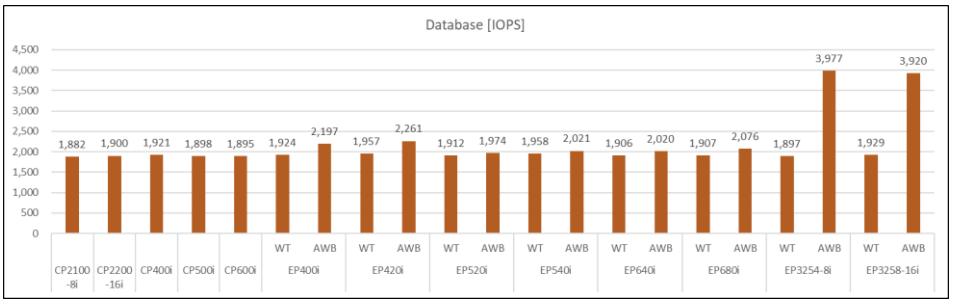


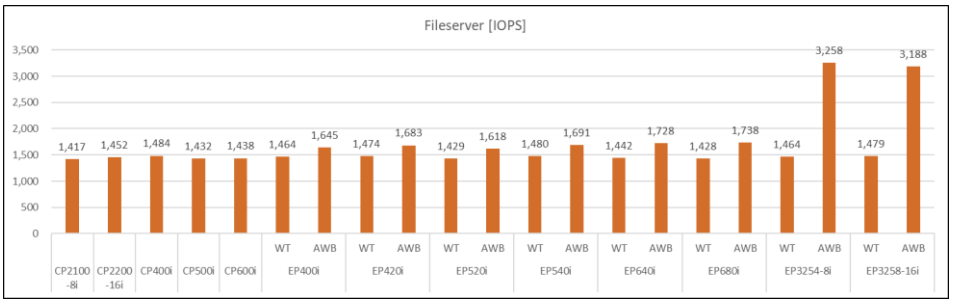


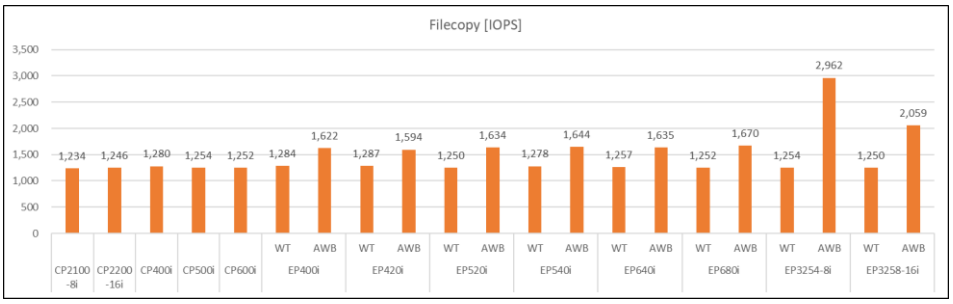


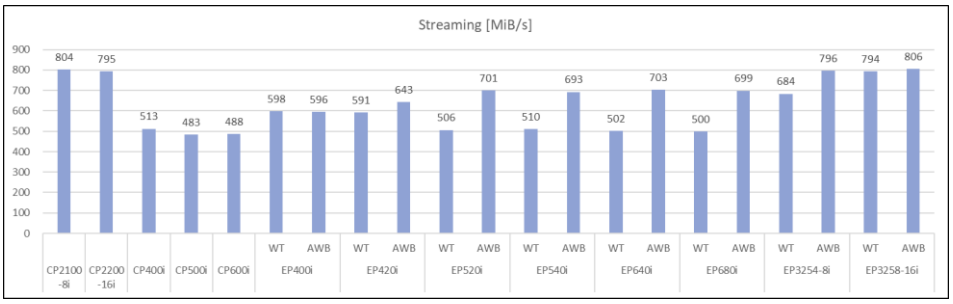


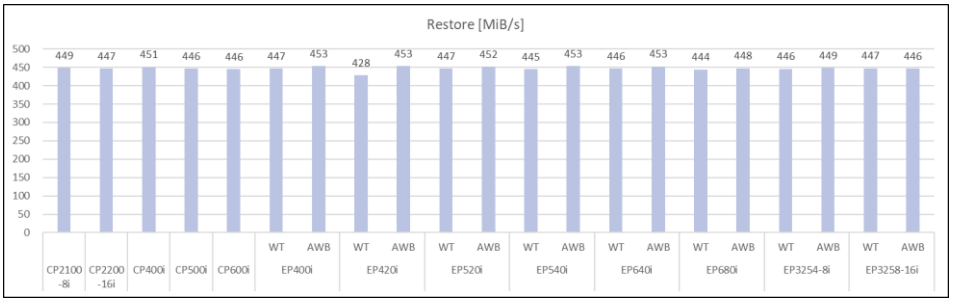


 

## [RAID5 4x drives]

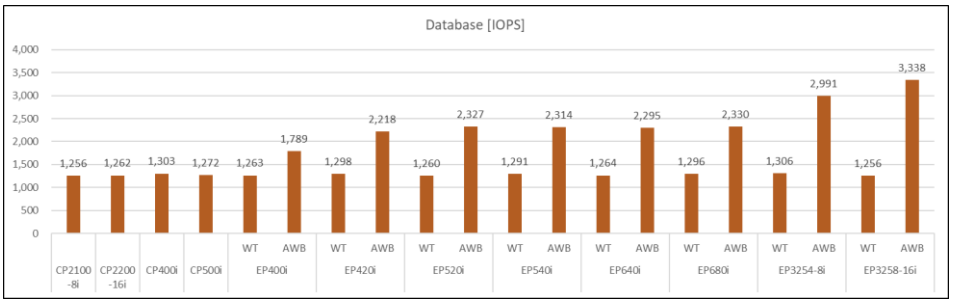


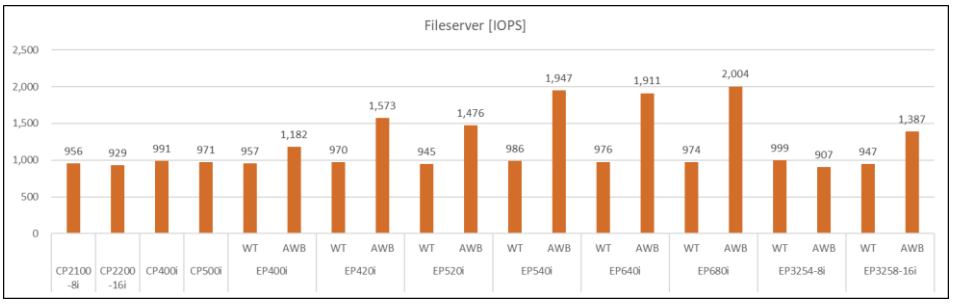


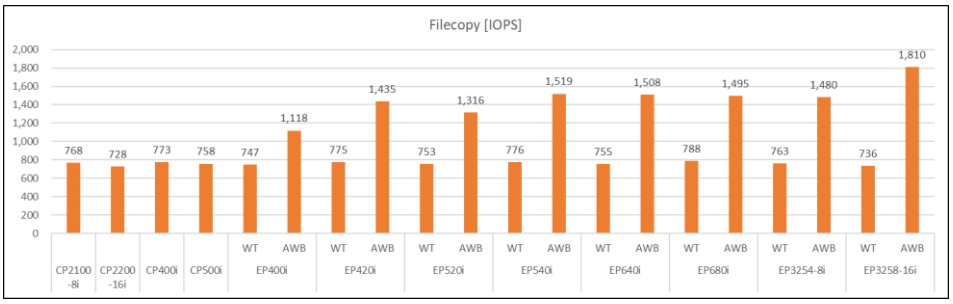


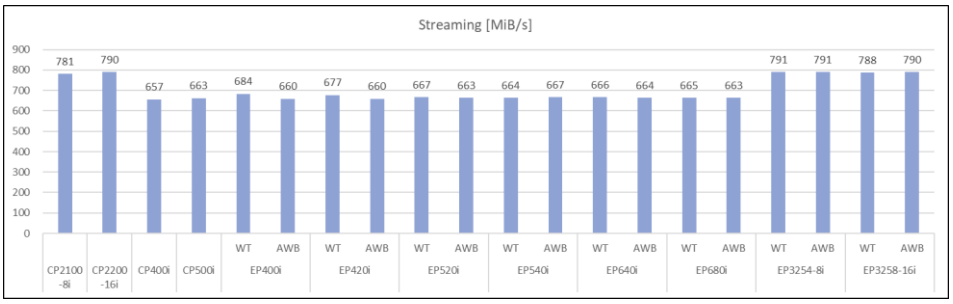


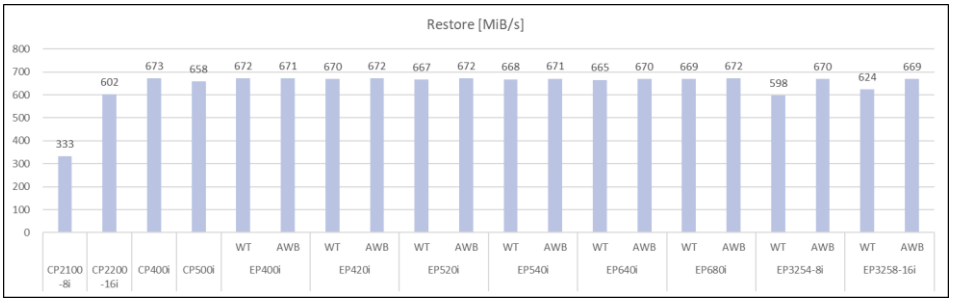


 

## *SAS-SSD Connection*

The figure below shows comparison data for a single drive and four drives at each RAID level.

With PRAID EP4x0i/5x0i/6x0i, enabling cache will result in significant performance degradation for database profiles with small data sizes. Other profiles improve by only about 10%.

With PRAID EP325x-xxi, write-through and write-back are similar performance.

We recommend using a write-through setting for SSD connections.

[RECOMMENDATION]

For SSD connections, select “Fast Path Optimum” for Cache Settings when creating logical drives using ServerView RAID Manager. A logical volume is created with the following RAID settings.

| Controller Cache policy | | Cache Mode | Disk Cache |
|---|---|---|---|
| Read | Write | | |
| No Read Ahead | Write Through | Direct | Enable |

[Single drive]

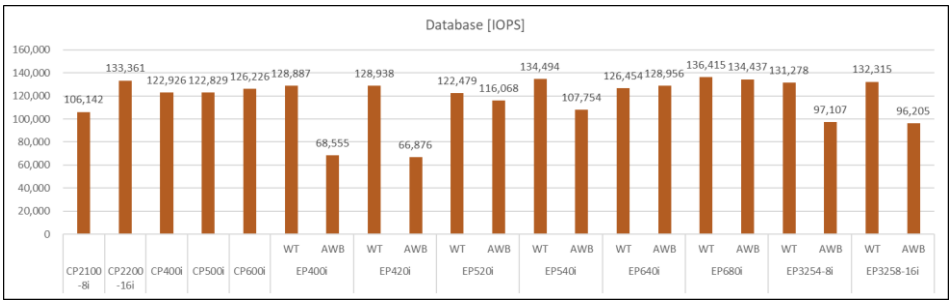


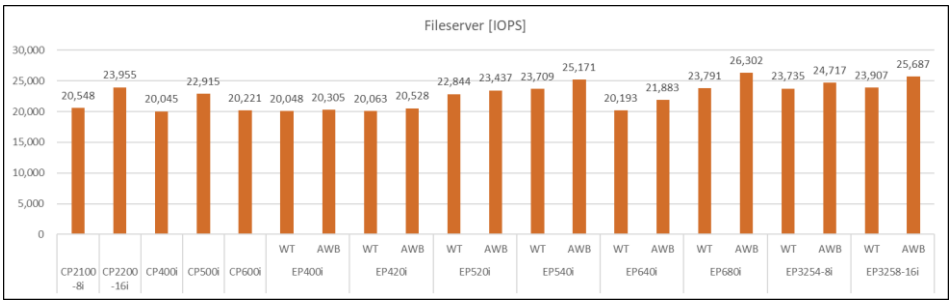


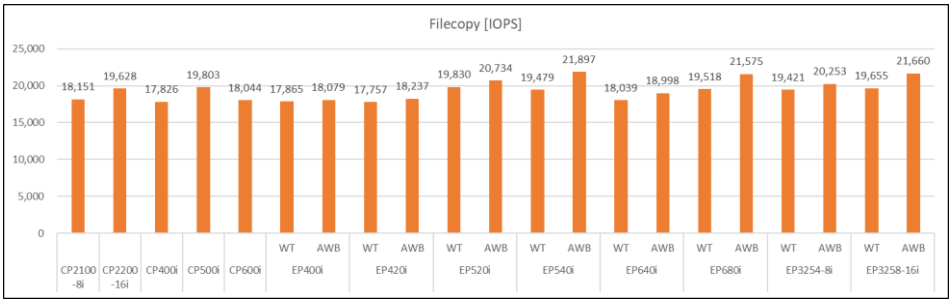


 

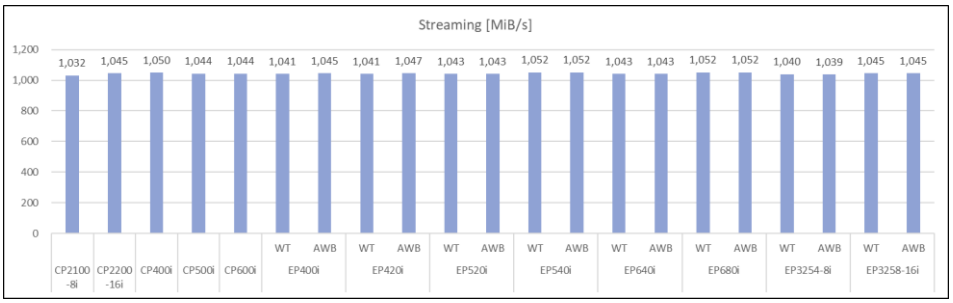


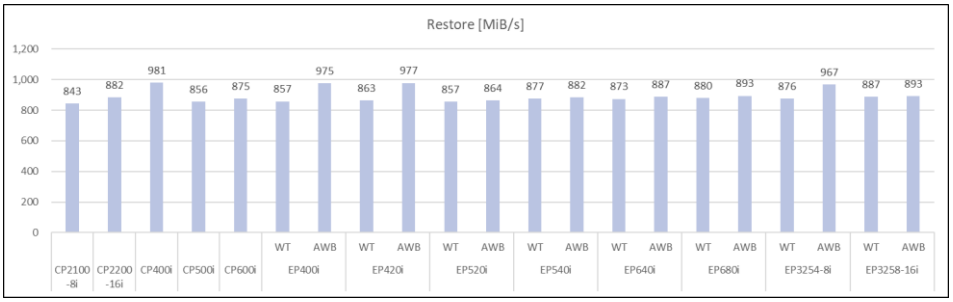


## [RAID0 4x drives]

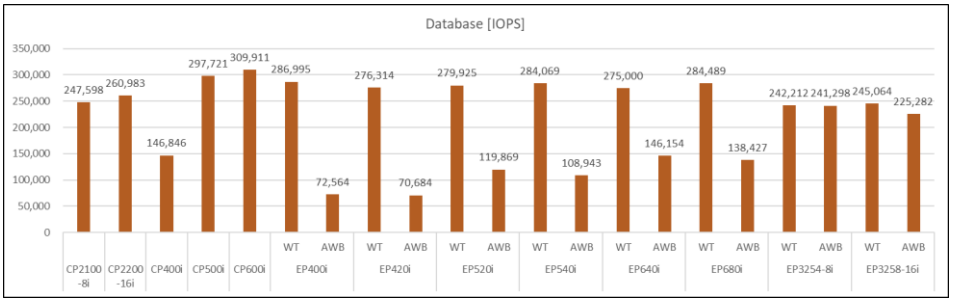


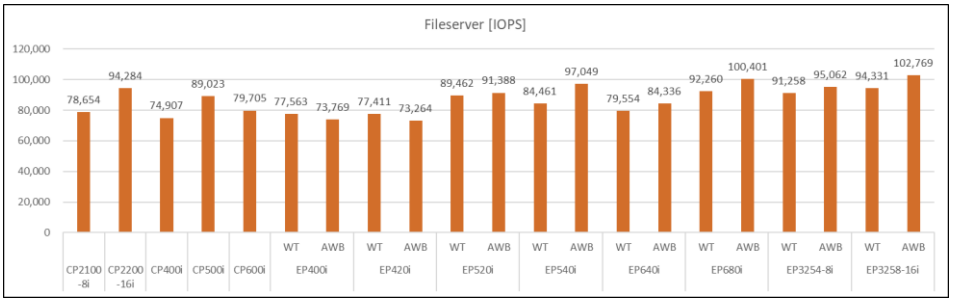


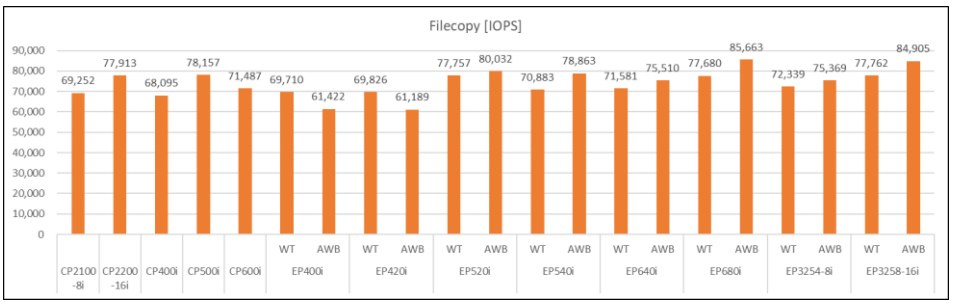


 

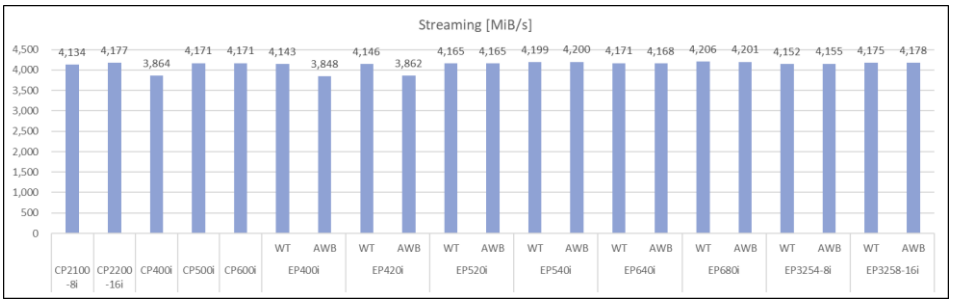


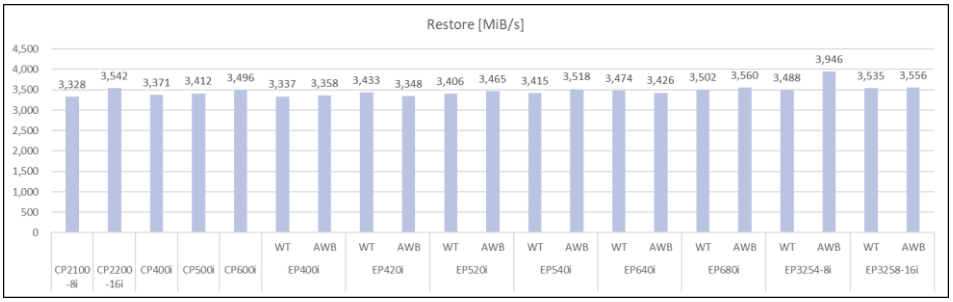


## [RAID10 4x drives]

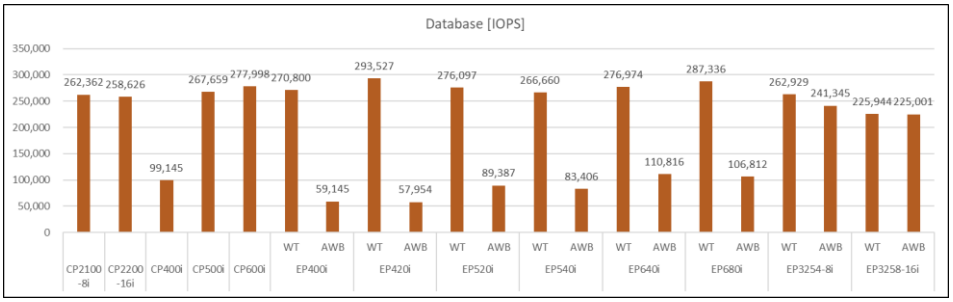


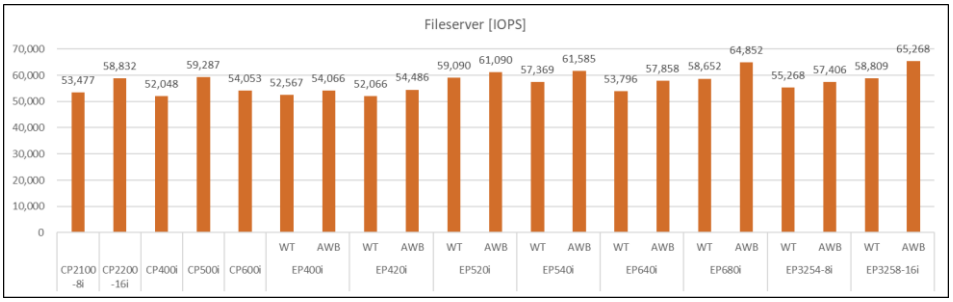


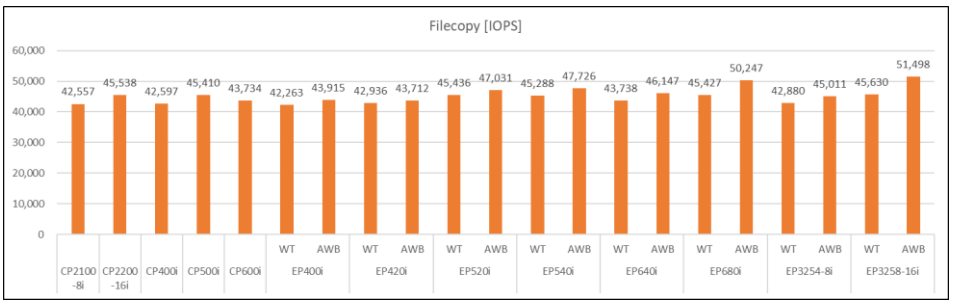


 

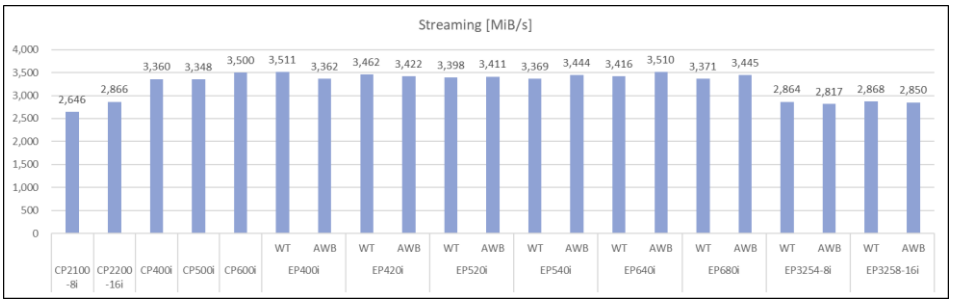


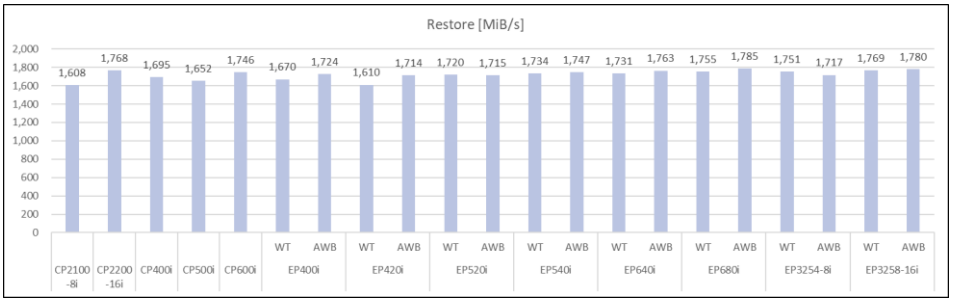


## [RAID5 4x drives]

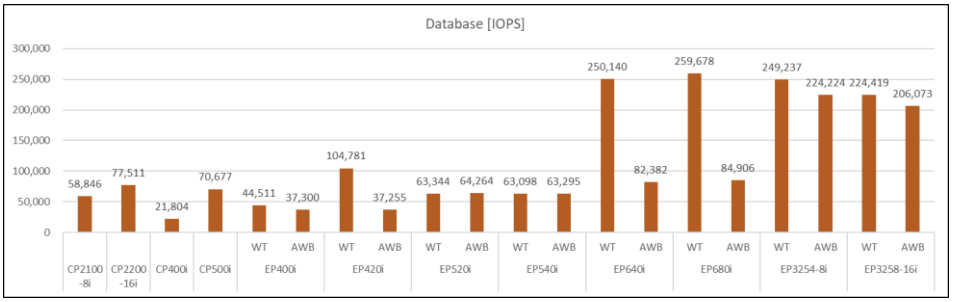


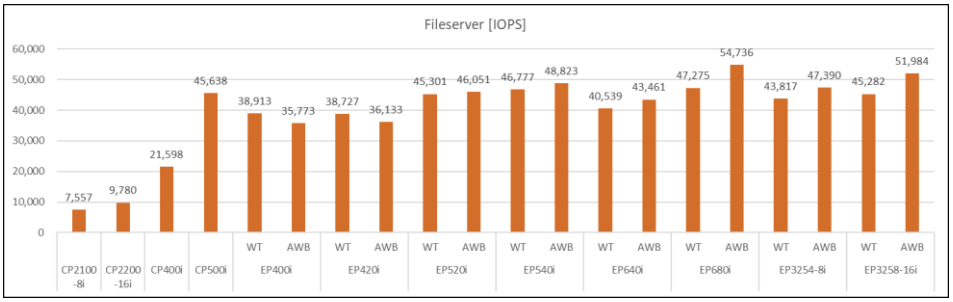


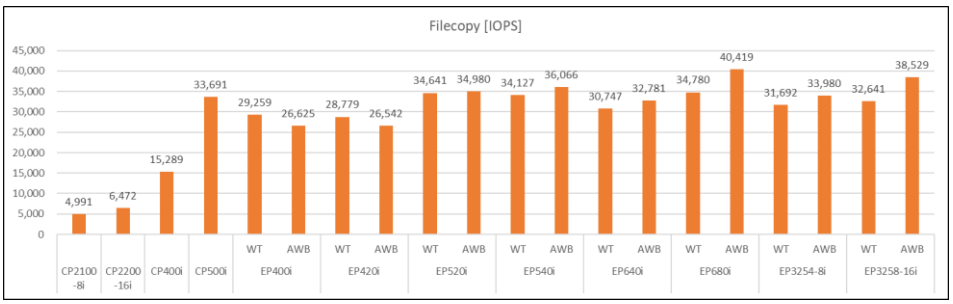


 

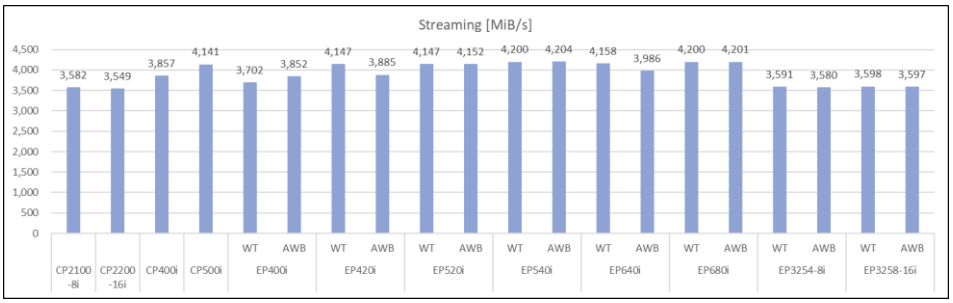


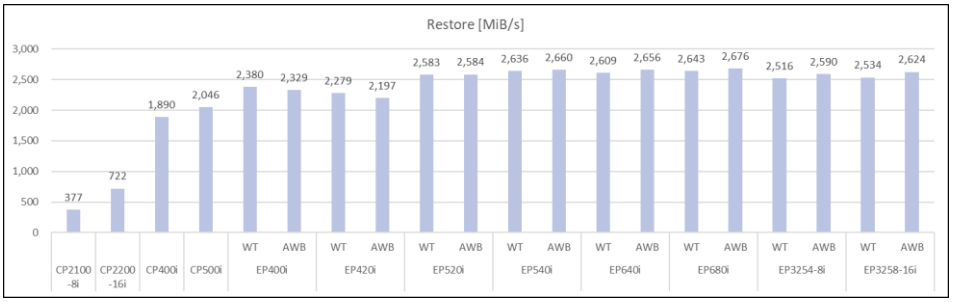


 

# Improved Transaction Rate with Multiplexing

The measurement results so far are the results measured with the number of worker threads as one in the IOMeter used as a measurement tool. In the case of processing with a small data size such as a database profile, the latency of the transfer processing becomes a bottleneck in SAS-SSD and NVMe drives, and the transaction rate has peaked. Increasing the number of worker threads can improve transaction rates by multiplexing the transfer process and increasing the load to the controller.

To see the processing power of the controller, the transaction rate was measured when the number of worker threads were increased in the database profile.

## *PRAID EP680i*

With a single drive, the transaction rate saturated with a small number of worker threads and stayed around 300k IOPS even with increasing the number of worker threads. With 4 drives, increasing the number of worker threads resulted in transaction rates exceeding 1,000k IOPS.

For RAID0

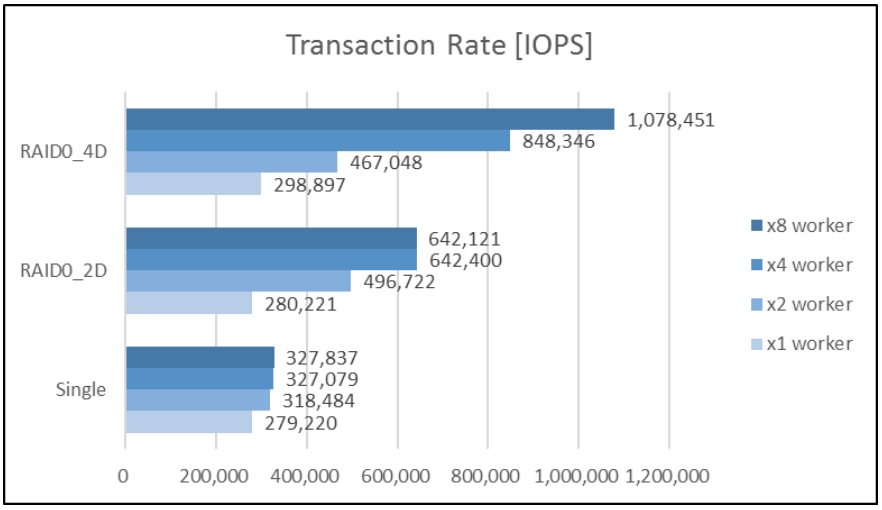


For RAID1/10

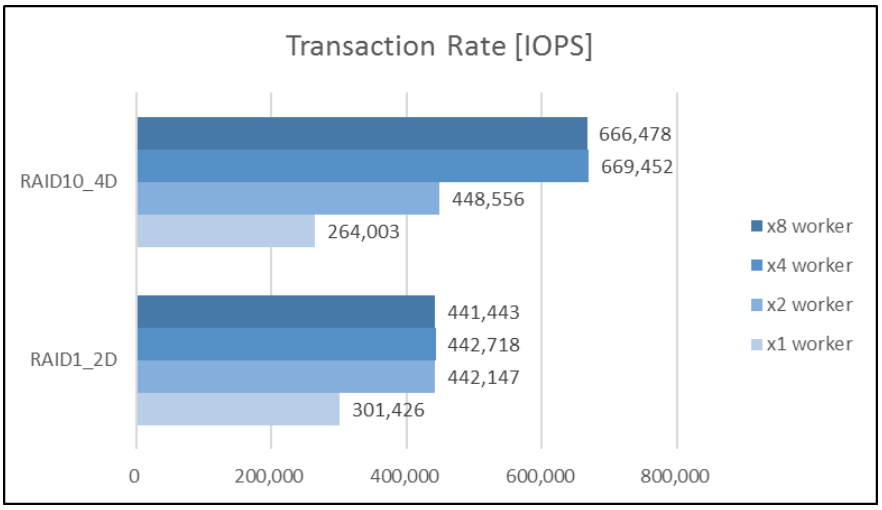


 

# Conclusion

The PRIMERGY and PRIMEQUEST servers support various RAID controllers.

The onboard controller can build a RAID 0/1/10 storage system without adding options by connecting the storage drive to the chipset's SATA interface or the processor's PCIe bus.

The PRAID CP500i/EP540i/EP680i installed in the PCIe slot supports the commonly used RAID level 0/1/10/5/50, and has the following functions for stable operation of the storage system.

- High reliability functions such as automatic drive media error repair, defective block replacement function, background media check function, RAID integrity assurance function, and hot spare are supported.
- Online capacity expansion (OCE) and online RAID level migration (RLM) allow you to change the array configuration while the system is running.
- By linking with the iRMC installed in the server itself, it is possible to monitor the status of storage drives and logical drives.

You can also monitor the status of logical drives and controllers from the operating system by installing ServerView RAID Manager.

The demand for storage is increasing more and more in order to meet the increasing amount of information year by year due to the spread of the Internet and cloud services.

HDDs are increasing in capacity due to various technological developments, and it is expected that they will continue to play an important role in storing large volumes of data. SSDs and NVMe drives are expected to be increasingly adopted in fields where high speed is required. For this reason, when designing a storage system, it is important to select the drive according to the application, and in some cases, it is effective to operate in a mixed manner. Also, when connecting a large number of SSD/NVMe drives to the controller, it is necessary to build a storage system considering the CPU side interface, which was not a problem with HDD. Consider the specifications of each RAID controller and build a storage system suitable for your application.

 

# Appendix 1: Results List

The measurement results are shown below.

## *HDD connection*

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| Intel VROC SATA RAID Controller | Single | 1 | 223 | 180 | 175 | 133 | 133 |
| | RAID0 | 2 | 439 | 354 | 345 | 264 | 264 |
| | | 4 | 864 | 701 | 682 | 525 | 523 |
| | | 8 | 1,658 | 1,360 | 1,327 | 1,028 | 1,041 |
| | RAID1/10 | 2 | 331 | 259 | 226 | 183 | 133 |
| | | 4 | 658 | 520 | 452 | 263 | 264 |
| Embedded MegaRAID | Single | 1 | 189 | 156 | 154 | 133 | 133 |
| | RAID0 | 2 | 319 | 270 | 269 | 266 | 246 |
| | | 4 | 524 | 455 | 464 | 532 | 310 |
| | | 8 | 831 | 726 | 762 | 1,049 | 398 |
| | RAID1/10 | 2 | 271 | 213 | 187 | 135 | 127 |
| | | 4 | 434 | 354 | 322 | 266 | 176 |
| | | 8 | 591 | 508 | 503 | 532 | 203 |
| PSAS CP2100-8i | Single | 1 | 586 | 448 | 439 | 232 | 230 |
| | RAID0 | 2 | 1,227 | 946 | 933 | 462 | 458 |
| | | 4 | 2,606 | 1,951 | 1,978 | 878 | 896 |
| | | 8 | 5,405 | 4,150 | 4,035 | 1,751 | 1,783 |
| | | 16 | 11,356 | 8,690 | 8,279 | 3,450 | 3,501 |
| | RAID1/10 | 2 | 893 | 666 | 574 | 434 | 230 |
| | | 4 | 1,882 | 1,417 | 1,234 | 804 | 449 |
| | | 8 | 4,029 | 2,959 | 2,589 | 1,416 | 897 |
| | | 16 | 8,738 | 6,522 | 5,686 | 2,290 | 1,744 |
| | RAID5 | 3 | 1,099 | 854 | 697 | 580 | 337 |
| | | 4 | 1,256 | 956 | 768 | 781 | 333 |
| | | 8 | 2,695 | 1,998 | 1,553 | 1,645 | 278 |
| | | 16 | 5,642 | 4,261 | 3,358 | 3,341 | 212 |
| PSAS CP2200-16i | Single | 1 | 601 | 461 | 446 | 229 | 230 |
| | RAID0 | 2 | 1,207 | 926 | 929 | 459 | 454 |
| | | 4 | 2,559 | 1,867 | 1,870 | 877 | 893 |
| | | 8 | 5,248 | 3,759 | 3,851 | 1,755 | 1,778 |
| | | 16 | 10,985 | 8,267 | 8,306 | 3,473 | 3,484 |
| | RAID1/10 | 2 | 907 | 691 | 603 | 406 | 229 |
| | | 4 | 1,900 | 1,452 | 1,246 | 795 | 447 |
| | | 8 | 3,965 | 2,993 | 2,546 | 1,383 | 869 |
| | | 16 | 8,487 | 6,336 | 5,407 | 2,060 | 824 |
| | RAID5 | 3 | 1,111 | 831 | 689 | 569 | 448 |
| | | 4 | 1,262 | 929 | 728 | 790 | 602 |
| | | 8 | 2,554 | 1,927 | 1,496 | 1,658 | 421 |
| | | 16 | 5,502 | 4,184 | 3,299 | 3,364 | 374 |
| PRAID CP400i | Single | 1 | 624 | 480 | 469 | 232 | 230 |
| | RAID0 | 2 | 1,286 | 984 | 957 | 463 | 456 |
| | | 4 | 2,655 | 1,984 | 1,926 | 875 | 886 |
| | | 8 | 5,172 | 3,900 | 3,840 | 1,726 | 1,535 |
| | | 16 | 9,362 | 7,851 | 7,971 | 3,422 | 1,540 |
| | RAID1/10 | 2 | 936 | 707 | 610 | 368 | 229 |
| | | 4 | 1,926 | 1,444 | 1,250 | 457 | 448 |
| | | 8 | 3,823 | 2,879 | 2,445 | 898 | 877 |
| | | 16 | 7,638 | 5,983 | 5,128 | 1,780 | 1,304 |
| | RAID5 | 3 | 1,154 | 877 | 718 | 454 | 446 |
| | | 4 | 1,283 | 959 | 737 | 661 | 605 |
| | | 8 | 2,492 | 1,909 | 1,465 | 1,537 | 954 |
| | | 16 | 4,805 | 3,880 | 3,103 | 3,239 | 976 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID CP500i | Single | 1 | 643 | 499 | 477 | 231 | 230 |
| | RAID0 | 2 | 1,294 | 996 | 953 | 463 | 456 |
| | | 4 | 2,664 | 1,997 | 1,900 | 908 | 893 |
| | | 8 | 5,458 | 4,009 | 4,005 | 1,739 | 1,767 |
| | | 16 | 10,914 | 8,649 | 8,542 | 3,454 | 3,424 |
| | RAID1/10 | 2 | 931 | 689 | 589 | 249 | 229 |
| | | 4 | 1,930 | 1,473 | 1,233 | 491 | 448 |
| | | 8 | 3,995 | 3,038 | 2,560 | 936 | 889 |
| | | 16 | 8,473 | 6,386 | 5,483 | 1,838 | 1,730 |
| | RAID5 | 3 | 1,144 | 880 | 697 | 454 | 442 |
| | | 4 | 1,320 | 1,002 | 763 | 682 | 672 |
| | | 8 | 2,677 | 2,021 | 1,591 | 1,537 | 1,528 |
| | | 16 | 5,749 | 4,413 | 3,403 | 3,249 | 883 |
| PRAID CP600i | Single | 1 | 651 | 504 | 495 | 232 | 229 |
| | RAID0 | 2 | 1,364 | 1,008 | 996 | 460 | 453 |
| | | 4 | 2,831 | 2,050 | 1,993 | 880 | 884 |
| | | 8 | 5,879 | 4,164 | 4,211 | 1,756 | 1,757 |
| | | 16 | 12,042 | 9,551 | 9,275 | 3,474 | 3,484 |
| | RAID1/10 | 2 | 972 | 726 | 627 | 249 | 226 |
| | | 4 | 2,044 | 1,525 | 1,359 | 488 | 445 |
| | | 8 | 4,322 | 3,214 | 2,766 | 927 | 882 |
| | | 16 | 8,777 | 6,602 | 5,641 | 1,838 | 1,754 |
| PRAID EP400i | Single | 1 | 611 | 464 | 453 | 232 | 230 |
| | RAID0 | 2 | 1,246 | 924 | 916 | 463 | 459 |
| | | 4 | 2,584 | 1,889 | 1,869 | 880 | 895 |
| | | 8 | 5,153 | 3,768 | 3,781 | 1,740 | 1,765 |
| | | 16 | 11,373 | 8,177 | 8,312 | 3,464 | 3,250 |
| | RAID1/10 | 2 | 946 | 692 | 598 | 377 | 229 |
| | | 4 | 1,944 | 1,497 | 1,290 | 689 | 446 |
| | | 8 | 4,022 | 3,042 | 2,554 | 1,329 | 873 |
| | | 16 | 8,270 | 6,236 | 5,343 | 2,383 | 1,661 |
| | RAID5 | 3 | 1,136 | 855 | 713 | 454 | 443 |
| | | 4 | 1,284 | 976 | 740 | 666 | 666 |
| | | 8 | 2,613 | 1,902 | 1,515 | 1,533 | 1,510 |
| | | 16 | 5,541 | 3,992 | 3,212 | 3,246 | 1,373 |
| PRAID EP420i | Single | 1 | 619 | 472 | 459 | 232 | 230 |
| | RAID0 | 2 | 1,276 | 960 | 942 | 463 | 459 |
| | | 4 | 2,583 | 1,929 | 1,885 | 881 | 892 |
| | | 8 | 5,291 | 3,860 | 3,851 | 1,738 | 1,764 |
| | | 16 | 11,315 | 8,152 | 8,205 | 3,475 | 3,270 |
| | RAID1/10 | 2 | 957 | 700 | 607 | 377 | 229 |
| | | 4 | 1,946 | 1,500 | 1,303 | 691 | 446 |
| | | 8 | 4,042 | 3,004 | 2,532 | 1,323 | 875 |
| | | 16 | 8,302 | 6,290 | 5,356 | 2,507 | 1,692 |
| | RAID5 | 3 | 1,154 | 877 | 709 | 454 | 441 |
| | | 4 | 1,279 | 952 | 743 | 662 | 665 |
| | | 8 | 2,575 | 1,850 | 1,500 | 1,529 | 1,515 |
| | | 16 | 5,589 | 4,090 | 3,176 | 3,265 | 1,325 |
| PRAID EP540i | Single | 1 | 634 | 498 | 475 | 231 | 229 |
| | RAID0 | 2 | 1,318 | 993 | 972 | 551 | 458 |
| | | 4 | 2,746 | 2,022 | 1,977 | 992 | 895 |
| | | 8 | 5,481 | 4,119 | 4,003 | 1,895 | 1,759 |
| | | 16 | 11,486 | 8,747 | 8,571 | 3,658 | 2,953 |
| | RAID1/10 | 2 | 945 | 685 | 597 | 467 | 229 |
| | | 4 | 1,914 | 1,472 | 1,263 | 704 | 448 |
| | | 8 | 4,103 | 3,042 | 2,526 | 1,375 | 880 |
| | | 16 | 8,349 | 6,390 | 5,426 | 2,600 | 1,712 |
| | RAID5 | 3 | 1,173 | 877 | 718 | 540 | 449 |
| | | 4 | 1,332 | 1,019 | 781 | 767 | 667 |
| | | 8 | 2,705 | 2,024 | 1,573 | 1,690 | 1,437 |
| | | 16 | 5,727 | 4,254 | 3,335 | 3,406 | 1,560 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID EP640i | Single | 1 | 617 | 476 | 477 | 225 | 223 |
| | RAID0 | 2 | 1,260 | 966 | 932 | 456 | 453 |
| | | 4 | 2,595 | 1,912 | 1,960 | 885 | 890 |
| | | 8 | 5,822 | 4,105 | 4,189 | 1,766 | 1,756 |
| | | 16 | 11,992 | 9,580 | 9,233 | 3,543 | 3,203 |
| | RAID1/10 | 2 | 930 | 685 | 586 | 248 | 228 |
| | | 4 | 1,910 | 1,464 | 1,228 | 481 | 445 |
| | | 8 | 3,988 | 3,005 | 2,550 | 1,257 | 881 |
| | | 16 | 8,182 | 6,257 | 5,094 | 2,342 | 1,719 |
| | RAID5 | 3 | 922 | 867 | 706 | 455 | 447 |
| | | 4 | 1,280 | 967 | 749 | 665 | 659 |
| | | 8 | 2,632 | 1,917 | 1,513 | 1,545 | 1,384 |
| | | 16 | 5,621 | 4,087 | 3,198 | 3,276 | 1,695 |
| PRAID EP680i | Single | 1 | 634 | 488 | 483 | 225 | 224 |
| | RAID0 | 2 | 1,288 | 979 | 947 | 457 | 453 |
| | | 4 | 2,702 | 2,009 | 2,005 | 881 | 889 |
| | | 8 | 5,823 | 4,212 | 4,200 | 1,766 | 1,759 |
| | | 16 | 12,008 | 9,590 | 9,339 | 3,499 | 3,233 |
| | RAID1/10 | 2 | 922 | 681 | 584 | 250 | 228 |
| | | 4 | 1,941 | 1,460 | 1,237 | 476 | 442 |
| | | 8 | 3,983 | 3,027 | 2,526 | 931 | 877 |
| | | 16 | 8,316 | 6,328 | 5,397 | 1,798 | 1,723 |
| | RAID5 | 3 | 1,148 | 865 | 693 | 454 | 447 |
| | | 4 | 1,335 | 961 | 768 | 666 | 661 |
| | | 8 | 2,816 | 2,028 | 1,638 | 1,538 | 1,412 |
| | | 16 | 5,977 | 4,743 | 3,592 | 3,265 | 1,627 |
| PRAID EP3254-8i | Single | 1 | 627 | 471 | 463 | 231 | 230 |
| | RAID0 | 2 | 1,293 | 970 | 955 | 460 | 457 |
| | | 4 | 2,674 | 1,954 | 1,915 | 872 | 894 |
| | | 8 | 5,480 | 3,914 | 4,024 | 1,742 | 1,765 |
| | | 16 | 11,444 | 8,564 | 8,506 | 3,461 | 3,488 |
| | RAID1/10 | 2 | 933 | 689 | 600 | 407 | 229 |
| | | 4 | 1,897 | 1,464 | 1,254 | 792 | 445 |
| | | 8 | 3,977 | 3,025 | 2,542 | 1,386 | 868 |
| | | 16 | 8,536 | 6,432 | 5,457 | 2,170 | 837 |
| | RAID5 | 3 | 1,149 | 881 | 710 | 341 | 228 |
| | | 4 | 1,306 | 999 | 763 | 684 | 446 |
| | | 8 | 2,658 | 2,017 | 1,585 | 1,211 | 863 |
| | | 16 | 5,607 | 4,236 | 3,294 | 2,011 | 809 |
| PRAID EP3258-16i | Single | 1 | 616 | 466 | 457 | 231 | 230 |
| | RAID0 | 2 | 1,241 | 942 | 916 | 460 | 455 |
| | | 4 | 2,550 | 1,891 | 1,906 | 879 | 888 |
| | | 8 | 5,210 | 3,842 | 3,905 | 1,751 | 1,608 |
| | | 16 | 11,162 | 8,287 | 8,417 | 3,448 | 3,452 |
| | RAID1/10 | 2 | 917 | 697 | 609 | 410 | 229 |
| | | 4 | 1,929 | 1,479 | 1,250 | 794 | 447 |
| | | 8 | 3,989 | 2,990 | 2,570 | 1,397 | 868 |
| | | 16 | 8,547 | 6,326 | 5,406 | 2,165 | 790 |
| | RAID5 | 3 | 1,122 | 846 | 709 | 570 | 449 |
| | | 4 | 1,256 | 947 | 736 | 788 | 624 |
| | | 8 | 2,548 | 1,961 | 1,544 | 1,673 | 490 |
| | | 16 | 5,468 | 4,080 | 3,178 | 3,377 | 562 |

 

## SSD connection

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| Intel VROC SATA RAID Controller | Single | 1 | 24,559 | 3,584 | 3,711 | 321 | 289 |
| | RAID0 | 2 | 46,286 | 7,022 | 7,301 | 648 | 529 |
| | | 4 | 87,851 | 14,201 | 14,033 | 1,299 | 884 |
| | | 8 | 171,743 | 27,109 | 28,070 | 1,893 | 1,620 |
| | RAID1/10 | 2 | 35,837 | 5,278 | 4,604 | 643 | 274 |
| | | 4 | 69,666 | 10,288 | 9,026 | 1,269 | 432 |
| Embedded MegaRAID | Single | 1 | 19,774 | 3,673 | 3,723 | 338 | 293 |
| | RAID0 | 2 | 29,085 | 5,952 | 6,124 | 520 | 527 |
| | | 4 | 41,793 | 8,822 | 9,253 | 901 | 978 |
| | | 8 | 62,971 | 12,744 | 13,443 | 1,555 | 1,697 |
| | RAID1/10 | 2 | 24,876 | 4,576 | 4,483 | 515 | 280 |
| | | 4 | 34,974 | 6,955 | 7,097 | 657 | 507 |
| | | 8 | 46,005 | 9,633 | 9,833 | 945 | 839 |
| PSAS CP2100-8i | Single | 1 | 106,142 | 20,548 | 18,151 | 1,032 | 843 |
| | RAID0 | 2 | 211,739 | 40,103 | 35,944 | 2,065 | 1,660 |
| | | 4 | 247,598 | 78,654 | 69,252 | 4,134 | 3,328 |
| | | 8 | 248,123 | 140,794 | 131,356 | 6,805 | 6,310 |
| | | 16 | 267,896 | 146,798 | 140,439 | 6,779 | 6,841 |
| | RAID1/10 | 2 | 153,015 | 27,211 | 21,578 | 1,322 | 831 |
| | | 4 | 262,362 | 53,477 | 42,557 | 2,646 | 1,608 |
| | | 8 | 265,340 | 100,269 | 80,926 | 5,311 | 3,193 |
| | | 16 | 244,353 | 108,284 | 89,112 | 6,809 | 3,424 |
| | RAID5 | 3 | 75,398 | 9,670 | 6,398 | 2,508 | 348 |
| | | 4 | 58,846 | 7,557 | 4,991 | 3,582 | 377 |
| | | 8 | 58,883 | 7,543 | 4,978 | 6,780 | 419 |
| | | 16 | 58,847 | 7,532 | 4,977 | 6,802 | 439 |
| PSAS CP2200-16i | Single | 1 | 133,361 | 23,955 | 19,628 | 1,045 | 882 |
| | RAID0 | 2 | 231,689 | 47,440 | 39,002 | 2,087 | 1,767 |
| | | 4 | 260,983 | 94,284 | 77,913 | 4,177 | 3,542 |
| | | 8 | 267,609 | 176,488 | 140,193 | 8,335 | 6,893 |
| | | 16 | 265,864 | 200,781 | 201,535 | 12,280 | 11,214 |
| | RAID1/10 | 2 | 196,759 | 29,481 | 22,222 | 1,436 | 886 |
| | | 4 | 258,626 | 58,832 | 45,538 | 2,866 | 1,768 |
| | | 8 | 263,334 | 116,688 | 90,479 | 5,566 | 3,527 |
| | | 16 | 265,154 | 190,628 | 144,365 | 10,135 | 5,856 |
| | RAID5 | 3 | 111,232 | 13,968 | 9,271 | 2,507 | 658 |
| | | 4 | 77,511 | 9,780 | 6,472 | 3,549 | 722 |
| | | 8 | 77,349 | 9,774 | 6,433 | 7,655 | 808 |
| | | 16 | 77,419 | 9,732 | 6,444 | 11,211 | 849 |
| PRAID CP400i | Single | 1 | 122,118 | 20,051 | 17,830 | 1,049 | 884 |
| | RAID0 | 2 | 143,398 | 39,284 | 35,357 | 2,005 | 1,614 |
| | | 4 | 155,026 | 77,015 | 69,107 | 3,860 | 3,340 |
| | | 8 | 153,802 | 120,070 | 121,534 | 5,843 | 5,349 |
| | | 16 | 154,874 | 114,732 | 128,281 | 5,895 | 6,429 |
| | RAID1/10 | 2 | 95,680 | 26,845 | 21,776 | 1,371 | 978 |
| | | 4 | 96,290 | 51,923 | 42,789 | 3,362 | 1,695 |
| | | 8 | 92,579 | 90,898 | 75,328 | 5,855 | 2,708 |
| | | 16 | 91,177 | 89,551 | 76,886 | 5,896 | 3,243 |
| | RAID5 | 3 | 25,096 | 24,344 | 17,998 | 2,989 | 1,622 |
| | | 4 | 22,342 | 21,865 | 15,247 | 3,861 | 1,903 |
| | | 8 | 22,613 | 22,891 | 16,182 | 5,852 | 1,965 |
| | | 16 | 22,555 | 22,526 | 15,685 | 5,898 | 1,975 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID CP500i | Single | 1 | 127,441 | 20,193 | 18,013 | 1,043 | 872 |
| | RAID0 | 2 | 236,180 | 40,122 | 35,870 | 2,084 | 1,737 |
| | | 4 | 294,598 | 79,422 | 70,862 | 4,173 | 3,481 |
| | | 8 | 307,849 | 113,052 | 137,944 | 7,047 | 6,556 |
| | | 16 | 289,179 | 149,216 | 150,754 | 7,063 | 6,917 |
| | RAID1/10 | 2 | 184,089 | 27,054 | 21,889 | 1,858 | 874 |
| | | 4 | 272,832 | 53,751 | 43,484 | 2,831 | 1,682 |
| | | 8 | 257,766 | 105,449 | 83,908 | 6,677 | 2,312 |
| | | 16 | 281,533 | 117,265 | 95,395 | 7,063 | 2,315 |
| | RAID5 | 3 | 82,506 | 34,386 | 26,639 | 3,108 | 1,637 |
| | | 4 | 71,002 | 37,848 | 29,379 | 4,161 | 2,075 |
| | | 8 | 74,721 | 51,282 | 34,770 | 7,054 | 2,190 |
| | | 16 | 73,956 | 50,961 | 34,852 | 7,063 | 2,247 |
| PRAID CP600i | Single | 1 | 126,569 | 19,776 | 18,043 | 1,044 | 875 |
| | RAID0 | 2 | 251,923 | 40,170 | 35,989 | 2,086 | 1,748 |
| | | 4 | 303,771 | 79,844 | 71,450 | 4,168 | 3,493 |
| | | 8 | 280,896 | 113,198 | 137,371 | 8,232 | 6,864 |
| | | 16 | 280,555 | 120,737 | 145,069 | 8,257 | 8,020 |
| | RAID1/10 | 2 | 185,376 | 27,148 | 21,950 | 1,820 | 811 |
| | | 4 | 280,932 | 54,031 | 43,726 | 3,504 | 1,750 |
| | | 8 | 285,869 | 107,720 | 83,360 | 6,833 | 3,487 |
| | | 16 | 277,637 | 121,728 | 98,407 | 8,307 | 4,041 |
| PRAID EP400i | Single | 1 | 128,725 | 20,005 | 17,873 | 1,041 | 864 |
| | RAID0 | 2 | 252,206 | 39,460 | 35,527 | 2,078 | 1,722 |
| | | 4 | 305,275 | 76,511 | 69,897 | 4,133 | 3,445 |
| | | 8 | 307,500 | 120,491 | 119,757 | 5,857 | 5,644 |
| | | 16 | 307,810 | 122,169 | 129,201 | 5,860 | 6,555 |
| | RAID1/10 | 2 | 188,857 | 26,491 | 21,590 | 1,769 | 936 |
| | | 4 | 271,284 | 52,740 | 41,985 | 3,493 | 1,688 |
| | | 8 | 271,767 | 92,082 | 75,637 | 5,850 | 2,756 |
| | | 16 | 272,614 | 95,545 | 86,392 | 5,856 | 3,239 |
| | RAID5 | 3 | 113,817 | 33,836 | 26,423 | 3,099 | 1,515 |
| | | 4 | 94,395 | 38,376 | 29,125 | 4,147 | 2,390 |
| | | 8 | 149,508 | 52,037 | 34,451 | 5,854 | 2,955 |
| | | 16 | 151,980 | 52,526 | 34,675 | 5,879 | 3,111 |
| PRAID EP420i | Single | 1 | 128,363 | 20,031 | 17,700 | 1,041 | 862 |
| | RAID0 | 2 | 254,206 | 39,048 | 35,528 | 2,077 | 1,724 |
| | | 4 | 307,229 | 77,842 | 68,109 | 4,142 | 3,443 |
| | | 8 | 298,773 | 120,489 | 120,236 | 5,859 | 5,655 |
| | | 16 | 309,666 | 121,928 | 129,170 | 5,845 | 6,550 |
| | RAID1/10 | 2 | 187,322 | 26,516 | 21,590 | 1,769 | 974 |
| | | 4 | 299,454 | 52,869 | 43,016 | 3,500 | 1,697 |
| | | 8 | 301,048 | 92,102 | 76,371 | 5,843 | 2,748 |
| | | 16 | 301,008 | 96,570 | 87,582 | 5,858 | 3,245 |
| | RAID5 | 3 | 120,742 | 33,884 | 26,591 | 2,971 | 1,730 |
| | | 4 | 84,758 | 38,761 | 28,920 | 4,146 | 2,326 |
| | | 8 | 148,942 | 51,157 | 33,678 | 5,858 | 2,885 |
| | | 16 | 148,894 | 51,272 | 33,973 | 5,862 | 3,052 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID EP540i | Single | 1 | 134,799 | 23,406 | 19,475 | 1,052 | 876 |
| | RAID0 | 2 | 261,098 | 46,127 | 38,735 | 2,101 | 1,751 |
| | | 4 | 267,580 | 86,541 | 74,938 | 4,207 | 3,488 |
| | | 8 | 271,264 | 120,391 | 123,082 | 6,776 | 6,242 |
| | | 16 | 272,001 | 127,314 | 122,801 | 6,750 | 6,520 |
| | RAID1/10 | 2 | 148,268 | 28,746 | 22,583 | 2,104 | 890 |
| | | 4 | 230,526 | 56,893 | 44,772 | 2,718 | 1,765 |
| | | 8 | 248,323 | 107,550 | 86,067 | 6,650 | 3,456 |
| | | 16 | 247,414 | 123,476 | 104,012 | 6,797 | 4,101 |
| | RAID5 | 3 | 71,618 | 37,857 | 28,195 | 3,155 | 1,762 |
| | | 4 | 61,808 | 46,010 | 34,077 | 4,207 | 2,643 |
| | | 8 | 83,781 | 60,863 | 40,466 | 6,782 | 3,450 |
| | | 16 | 82,820 | 61,893 | 41,053 | 6,794 | 3,464 |
| PRAID EP640i | Single | 1 | 126,336 | 20,189 | 18,036 | 1,044 | 875 |
| | RAID0 | 2 | 250,703 | 40,189 | 35,838 | 2,086 | 1,751 |
| | | 4 | 259,432 | 79,095 | 70,970 | 4,171 | 3,376 |
| | | 8 | 297,621 | 144,283 | 133,724 | 8,287 | 6,183 |
| | | 16 | 265,413 | 178,524 | 160,640 | 8,325 | 7,948 |
| | RAID1/10 | 2 | 183,259 | 27,125 | 21,956 | 1,856 | 872 |
| | | 4 | 248,530 | 53,961 | 43,708 | 3,501 | 1,750 |
| | | 8 | 266,129 | 106,246 | 82,181 | 6,749 | 3,448 |
| | | 16 | 268,501 | 111,962 | 98,241 | 8,307 | 4,040 |
| | RAID5 | 3 | 176,639 | 34,554 | 27,336 | 3,110 | 1,760 |
| | | 4 | 248,777 | 40,044 | 29,438 | 4,159 | 2,625 |
| | | 8 | 256,640 | 78,776 | 56,602 | 8,271 | 4,782 |
| | | 16 | 277,408 | 85,733 | 57,420 | 8,325 | 5,241 |
| PRAID EP680i | Single | 1 | 135,667 | 23,904 | 19,608 | 1,052 | 883 |
| | RAID0 | 2 | 271,118 | 47,414 | 39,085 | 2,104 | 1,764 |
| | | 4 | 298,251 | 91,565 | 77,126 | 4,207 | 3,511 |
| | | 8 | 296,201 | 179,236 | 133,505 | 8,348 | 6,829 |
| | | 16 | 286,457 | 218,977 | 215,577 | 13,956 | 11,627 |
| | RAID1/10 | 2 | 201,979 | 29,507 | 22,842 | 2,100 | 882 |
| | | 4 | 279,795 | 58,407 | 45,483 | 3,454 | 1,764 |
| | | 8 | 264,444 | 111,607 | 88,669 | 6,818 | 3,423 |
| | | 16 | 285,461 | 193,389 | 176,378 | 13,239 | 6,668 |
| | RAID5 | 3 | 169,763 | 39,611 | 27,623 | 2,585 | 1,785 |
| | | 4 | 251,934 | 46,113 | 33,826 | 4,206 | 2,575 |
| | | 8 | 248,135 | 82,087 | 55,927 | 8,371 | 4,775 |
| | | 16 | 269,865 | 84,559 | 56,187 | 14,119 | 4,968 |
| PRAID EP3254-8i | Single | 1 | 131,278 | 23,735 | 19,421 | 1,040 | 876 |
| | RAID0 | 2 | 251,426 | 46,509 | 38,133 | 2,080 | 1,751 |
| | | 4 | 242,212 | 91,258 | 72,339 | 4,152 | 3,488 |
| | | 8 | 267,138 | 181,449 | 131,641 | 8,267 | 6,719 |
| | | 16 | 266,166 | 185,270 | 159,787 | 8,281 | 7,271 |
| | RAID1/10 | 2 | 195,861 | 28,980 | 22,454 | 1,424 | 879 |
| | | 4 | 262,929 | 55,268 | 42,880 | 2,864 | 1,751 |
| | | 8 | 241,880 | 96,504 | 74,609 | 5,492 | 3,462 |
| | | 16 | 241,652 | 116,309 | 85,756 | 8,309 | 3,605 |
| | RAID5 | 3 | 188,978 | 38,130 | 28,209 | 2,512 | 1,759 |
| | | 4 | 249,237 | 43,817 | 31,692 | 3,591 | 2,516 |
| | | 8 | 266,298 | 62,705 | 41,641 | 7,768 | 4,705 |
| | | 16 | 239,253 | 62,901 | 41,620 | 8,250 | 4,912 |
| PRAID EP3258-16i | Single | 1 | 132,315 | 23,907 | 19,655 | 1,045 | 887 |
| | RAID0 | 2 | 241,202 | 47,518 | 39,083 | 2,089 | 1,763 |
| | | 4 | 245,064 | 94,331 | 77,762 | 4,175 | 3,535 |
| | | 8 | 245,679 | 181,829 | 154,034 | 8,341 | 7,057 |
| | | 16 | 227,158 | 173,638 | 177,075 | 10,576 | 10,959 |
| | RAID1/10 | 2 | 194,359 | 29,538 | 22,860 | 1,415 | 885 |
| | | 4 | 225,944 | 58,809 | 45,630 | 2,868 | 1,769 |
| | | 8 | 227,944 | 116,315 | 90,307 | 5,728 | 3,528 |
| | | 16 | 225,791 | 173,480 | 172,057 | 10,209 | 6,326 |
| | RAID5 | 3 | 197,523 | 39,090 | 29,151 | 2,519 | 1,765 |
| | | 4 | 224,419 | 45,282 | 32,641 | 3,598 | 2,534 |
| | | 8 | 245,448 | 62,773 | 41,741 | 7,811 | 4,778 |
| | | 16 | 226,138 | 63,181 | 41,659 | 10,482 | 5,027 |

 

# NVMe connection

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| Intel VROC NVMe RAID Controller | Single | 1 | 208,059 | 65,632 | 80,569 | 2,845 | 2,820 |
| | RAID0 | 2 | 175,483 | 130,895 | 144,043 | 5,690 | 5,512 |
| | | 4 | 177,018 | 150,207 | 149,187 | 9,577 | 9,046 |
| | | 8 | 175,798 | 161,191 | 160,263 | 9,994 | 9,823 |
| | RAID1/10 | 2 | 174,566 | 117,297 | 82,509 | 5,690 | 2,759 |
| | | 4 | 194,805 | 156,594 | 151,426 | 10,842 | 5,250 |
| | RAID5 | 3 | 176,026 | 113,090 | 82,393 | 8,535 | 2,588 |
| | | 4 | 161,897 | 98,987 | 72,419 | 9,167 | 3,020 |
| | | 8 | 148,279 | 69,350 | 50,370 | 8,814 | 2,201 |
| Microsoft storage pool | Single | 1 | 193,457 | 48,715 | 46,109 | 3,189 | 2,785 |
| Intel PCIe Device Adapter | Simple | 2 | 181,967 | 94,626 | 90,032 | 6,409 | 5,446 |
| | | 4 | 189,475 | 127,781 | 122,901 | 8,311 | 6,573 |
| | | 8 | 183,145 | 123,954 | 120,481 | 8,497 | 6,652 |
| | Mirror | 2 | 148,669 | 63,698 | 55,501 | 3,206 | 2,709 |
| | | 4 | 144,915 | 96,264 | 83,021 | 8,726 | 4,076 |
| | | 8 | 140,606 | 94,734 | 83,503 | 8,441 | 3,988 |
| PRAID EP540i | Single | 1 | 271,172 | 49,053 | 48,154 | 3,190 | 2,513 |
| | RAID0 | 2 | 256,804 | 98,887 | 96,366 | 6,364 | 4,805 |
| | | 4 | 257,050 | 149,826 | 169,427 | 6,699 | 6,499 |
| | RAID1/10 | 2 | 254,186 | 68,303 | 58,103 | 6,379 | 2,451 |
| | | 4 | 253,374 | 126,197 | 96,298 | 6,711 | 3,601 |
| | RAID5 | 3 | 73,443 | 66,018 | 43,543 | 6,666 | 2,519 |
| | | 4 | 62,829 | 55,373 | 36,559 | 6,684 | 3,021 |
| PRAID EP680i | Single | 1 | 260,331 | 49,771 | 49,216 | 3,212 | 2,619 |
| | RAID0 | 2 | 275,019 | 79,960 | 78,514 | 6,424 | 5,146 |
| | | 4 | 263,631 | 158,357 | 142,111 | 10,399 | 8,848 |
| | RAID1/10 | 2 | 256,873 | 70,538 | 60,743 | 6,422 | 2,554 |
| | | 4 | 262,427 | 123,860 | 104,826 | 10,102 | 4,933 |
| | RAID5 | 3 | 100,375 | 74,812 | 59,372 | 8,938 | 3,418 |
| | | 4 | 85,242 | 73,768 | 49,380 | 10,278 | 4,123 |
| PSAS CP2200-16i | Single | 1 | 239,374 | 53,073 | 49,707 | 3,227 | 2,875 |
| | RAID0 | 2 | 267,199 | 100,469 | 91,445 | 6,130 | 5,460 |
| | | 4 | 245,692 | 179,203 | 178,266 | 10,967 | 8,878 |
| | RAID1/10 | 2 | 268,679 | 75,164 | 61,707 | 6,252 | 2,768 |
| | | 4 | 267,291 | 138,135 | 114,130 | 11,661 | 5,366 |
| | RAID5 | 3 | 97,766 | 12,584 | 8,345 | 9,682 | 607 |
| | | 4 | 68,358 | 8,755 | 5,798 | 11,241 | 684 |
| PRAID EP3258-16i | Single | 1 | 228,029 | 52,656 | 49,653 | 3,227 | 2,814 |
| | RAID0 | 2 | 249,600 | 104,722 | 97,041 | 6,455 | 5,622 |
| | | 4 | 250,712 | 183,398 | 184,792 | 11,519 | 10,693 |
| | RAID1/10 | 2 | 229,793 | 74,962 | 63,092 | 6,455 | 2,794 |
| | | 4 | 247,877 | 143,397 | 122,115 | 11,379 | 5,309 |
| | RAID5 | 3 | 243,595 | 81,650 | 54,939 | 9,682 | 3,810 |
| | | 4 | 245,053 | 58,208 | 38,684 | 10,847 | 4,193 |

 

# Appendix 2 : Cache effect measurement results

Detailed data of the measurement results made to check the effect of the cache.

## *HDD connection*

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PSAS CP2100-8i | Single | 1 | 586 | 448 | 439 | 232 | 230 |
| | RAID0 | 2 | 1,227 | 946 | 933 | 462 | 458 |
| | | 4 | 2,606 | 1,951 | 1,978 | 878 | 896 |
| | | 8 | 5,405 | 4,150 | 4,035 | 1,751 | 1,783 |
| | | 16 | 11,356 | 8,690 | 8,279 | 3,450 | 3,501 |
| | RAID1/10 | 2 | 893 | 666 | 574 | 434 | 230 |
| | | 4 | 1,882 | 1,417 | 1,234 | 804 | 449 |
| | | 8 | 4,029 | 2,959 | 2,589 | 1,416 | 897 |
| | | 16 | 8,738 | 6,522 | 5,686 | 2,290 | 1,744 |
| | RAID5 | 3 | 1,099 | 854 | 697 | 580 | 337 |
| | | 4 | 1,256 | 956 | 768 | 781 | 333 |
| | | 8 | 2,695 | 1,998 | 1,553 | 1,645 | 278 |
| | | 16 | 5,642 | 4,261 | 3,358 | 3,341 | 212 |
| PSAS CP2200-16i | Single | 1 | 601 | 461 | 446 | 229 | 230 |
| | RAID0 | 2 | 1,207 | 926 | 929 | 459 | 454 |
| | | 4 | 2,559 | 1,867 | 1,870 | 877 | 893 |
| | | 8 | 5,248 | 3,759 | 3,851 | 1,755 | 1,778 |
| | | 16 | 10,985 | 8,267 | 8,306 | 3,473 | 3,484 |
| | RAID1/10 | 2 | 907 | 691 | 603 | 406 | 229 |
| | | 4 | 1,900 | 1,452 | 1,246 | 795 | 447 |
| | | 8 | 3,965 | 2,993 | 2,546 | 1,383 | 869 |
| | | 16 | 8,487 | 6,336 | 5,407 | 2,060 | 824 |
| | RAID5 | 3 | 1,111 | 831 | 689 | 569 | 448 |
| | | 4 | 1,262 | 929 | 728 | 790 | 602 |
| | | 8 | 2,554 | 1,927 | 1,496 | 1,658 | 421 |
| | | 16 | 5,502 | 4,184 | 3,299 | 3,364 | 374 |
| PRAID CP400i | Single | 1 | 622 | 474 | 468 | 232 | 230 |
| | RAID0 | 2 | 1,284 | 982 | 972 | 463 | 457 |
| | | 4 | 2,649 | 1,985 | 1,983 | 880 | 896 |
| | | 8 | 5,437 | 4,021 | 4,039 | 1,750 | 1,757 |
| | | 16 | 11,160 | 8,365 | 8,331 | 3,463 | 2,887 |
| | RAID1/10 | 2 | 951 | 706 | 609 | 367 | 230 |
| | | 4 | 1,921 | 1,484 | 1,280 | 513 | 451 |
| | | 8 | 4,043 | 3,006 | 2,646 | 904 | 878 |
| | | 16 | 8,413 | 6,334 | 5,574 | 1,777 | 1,730 |
| | RAID5 | 3 | 1,155 | 880 | 727 | 455 | 450 |
| | | 4 | 1,303 | 991 | 773 | 657 | 673 |
| | | 8 | 2,707 | 2,008 | 1,606 | 1,533 | 1,190 |
| | | 16 | 5,614 | 4,231 | 3,344 | 3,226 | 1,193 |
| PRAID CP500i | Single | 1 | 604 | 470 | 468 | 232 | 228 |
| | RAID0 | 2 | 1,254 | 955 | 943 | 463 | 456 |
| | | 4 | 2,618 | 1,923 | 1,932 | 887 | 894 |
| | | 8 | 5,357 | 3,935 | 3,933 | 1,767 | 1,764 |
| | | 16 | 11,389 | 8,327 | 8,342 | 3,499 | 3,473 |
| | RAID1/10 | 2 | 937 | 675 | 588 | 250 | 229 |
| | | 4 | 1,898 | 1,432 | 1,254 | 483 | 446 |
| | | 8 | 3,979 | 2,976 | 2,583 | 935 | 887 |
| | | 16 | 8,416 | 6,284 | 5,500 | 1,838 | 1,720 |
| | RAID5 | 3 | 1,130 | 845 | 697 | 454 | 445 |
| | | 4 | 1,272 | 971 | 758 | 663 | 658 |
| | | 8 | 2,595 | 1,930 | 1,535 | 1,536 | 1,500 |
| | | 16 | 5,681 | 4,113 | 3,286 | 3,274 | 941 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID CP600i | Single | 1 | 602 | 470 | 462 | 232 | 228 |
| | RAID0 | 2 | 1,249 | 944 | 948 | 460 | 452 |
| | | 4 | 2,600 | 1,902 | 1,937 | 883 | 888 |
| | | 8 | 5,332 | 3,890 | 3,989 | 1,753 | 1,767 |
| | | 16 | 11,369 | 8,425 | 8,304 | 3,484 | 3,472 |
| | RAID1/10 | 2 | 934 | 676 | 588 | 249 | 227 |
| | | 4 | 1,895 | 1,438 | 1,252 | 488 | 446 |
| | | 8 | 3,970 | 2,977 | 2,617 | 939 | 885 |
| | | 16 | 8,361 | 6,329 | 5,523 | 1,835 | 1,743 |
| PRAID EP400i WT: Write Through | Single | 1 | 609 | 459 | 458 | 232 | 229 |
| | RAID0 | 2 | 1,259 | 938 | 955 | 482 | 457 |
| | | 4 | 2,569 | 1,877 | 1,931 | 879 | 884 |
| | | 8 | 5,332 | 3,810 | 3,916 | 1,758 | 1,682 |
| | | 16 | 11,392 | 8,315 | 8,133 | 3,475 | 2,954 |
| | RAID1/10 | 2 | 942 | 705 | 606 | 373 | 229 |
| | | 4 | 1,924 | 1,464 | 1,284 | 598 | 447 |
| | | 8 | 4,049 | 3,048 | 2,648 | 1,179 | 874 |
| | | 16 | 8,369 | 6,365 | 5,487 | 1,936 | 1,663 |
| | RAID5 | 3 | 1,111 | 862 | 707 | 453 | 448 |
| | | 4 | 1,263 | 957 | 747 | 684 | 672 |
| | | 8 | 2,506 | 1,909 | 1,517 | 1,531 | 1,502 |
| | | 16 | 5,678 | 4,034 | 3,201 | 3,275 | 1,302 |
| PRAID EP400i AWB: Always Write Back | Single | 1 | 744 | 563 | 571 | 232 | 229 |
| | RAID0 | 2 | 1,351 | 979 | 1,109 | 461 | 457 |
| | | 4 | 2,758 | 1,941 | 2,012 | 892 | 894 |
| | | 8 | 5,682 | 3,818 | 4,053 | 1,764 | 1,784 |
| | | 16 | 11,430 | 8,144 | 8,663 | 3,476 | 3,553 |
| | RAID1/10 | 2 | 1,196 | 943 | 876 | 358 | 217 |
| | | 4 | 2,197 | 1,645 | 1,622 | 596 | 453 |
| | | 8 | 4,473 | 3,332 | 3,147 | 1,165 | 898 |
| | | 16 | 9,044 | 7,035 | 6,497 | 2,042 | 1,793 |
| | RAID5 | 3 | 2,063 | 1,067 | 891 | 452 | 447 |
| | | 4 | 1,789 | 1,182 | 1,118 | 660 | 671 |
| | | 8 | 2,996 | 2,299 | 2,011 | 1,552 | 1,568 |
| | | 16 | 6,575 | 4,593 | 3,928 | 3,273 | 3,278 |
| PRAID EP420i WT: Write Through | Single | 1 | 614 | 468 | 467 | 231 | 229 |
| | RAID0 | 2 | 1,296 | 969 | 973 | 507 | 456 |
| | | 4 | 2,672 | 1,957 | 1,962 | 876 | 884 |
| | | 8 | 5,424 | 3,928 | 3,865 | 1,775 | 1,692 |
| | | 16 | 11,604 | 8,439 | 8,325 | 3,486 | 2,927 |
| | RAID1/10 | 2 | 943 | 701 | 616 | 374 | 228 |
| | | 4 | 1,957 | 1,474 | 1,287 | 591 | 428 |
| | | 8 | 4,055 | 3,053 | 2,671 | 1,157 | 874 |
| | | 16 | 8,552 | 6,421 | 5,560 | 1,937 | 1,686 |
| | RAID5 | 3 | 1,138 | 872 | 722 | 452 | 448 |
| | | 4 | 1,298 | 970 | 775 | 677 | 670 |
| | | 8 | 2,656 | 1,909 | 1,511 | 1,556 | 1,458 |
| | | 16 | 5,750 | 4,038 | 3,181 | 3,258 | 1,368 |
| PRAID EP420i AWB: Always Write Back | Single | 1 | 881 | 586 | 704 | 231 | 229 |
| | RAID0 | 2 | 1,458 | 1,027 | 1,286 | 461 | 459 |
| | | 4 | 2,839 | 2,017 | 2,119 | 918 | 894 |
| | | 8 | 5,846 | 3,854 | 4,117 | 1,777 | 1,792 |
| | | 16 | 11,153 | 8,234 | 8,562 | 3,478 | 3,561 |
| | RAID1/10 | 2 | 1,410 | 1,001 | 1,049 | 376 | 233 |
| | | 4 | 2,261 | 1,683 | 1,594 | 643 | 453 |
| | | 8 | 4,574 | 3,376 | 3,178 | 1,272 | 899 |
| | | 16 | 9,187 | 7,061 | 6,682 | 1,962 | 1,805 |
| | RAID5 | 3 | 2,063 | 1,370 | 1,297 | 453 | 448 |
| | | 4 | 2,218 | 1,573 | 1,435 | 660 | 672 |
| | | 8 | 3,111 | 2,604 | 1,951 | 1,548 | 1,566 |
| | | 16 | 6,826 | 4,681 | 4,004 | 3,276 | 3,216 |
| PRAID EP520i WT: Write Through | Single | 1 | 603 | 464 | 455 | 232 | 229 |
| | RAID0 | 2 | 1,253 | 953 | 952 | 498 | 455 |
| | | 4 | 2,566 | 1,930 | 1,929 | 921 | 888 |
| | | 8 | 5,350 | 3,842 | 3,922 | 1,799 | 1,751 |
| | | 16 | 11,341 | 8,255 | 8,067 | 3,528 | 3,255 |
| | RAID1/10 | 2 | 933 | 688 | 591 | 250 | 229 |
| | | 4 | 1,912 | 1,429 | 1,250 | 506 | 447 |
| | | 8 | 3,983 | 2,977 | 2,593 | 968 | 885 |
| | | 16 | 8,344 | 6,168 | 5,451 | 1,857 | 1,731 |
| | RAID5 | 3 | 1,106 | 836 | 699 | 455 | 447 |
| | | 4 | 1,260 | 945 | 753 | 667 | 667 |
| | | 8 | 2,608 | 1,895 | 1,515 | 1,573 | 1,384 |
| | | 16 | 5,664 | 3,969 | 3,092 | 3,306 | 1,652 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID EP520i | Single | 1 | 687 | 581 | 624 | 232 | 229 |
| AWB: | RAID0 | 2 | 1,265 | 970 | 1,249 | 461 | 457 |
| Always Write Back | | 4 | 2,840 | 1,995 | 2,080 | 885 | 898 |
| | | 8 | 5,810 | 3,939 | 4,251 | 1,796 | 1,788 |
| | | 16 | 10,908 | 8,350 | 8,486 | 3,546 | 3,562 |
| | RAID1/10 | 2 | 1,046 | 888 | 986 | 377 | 229 |
| | | 4 | 1,974 | 1,618 | 1,634 | 701 | 452 |
| | | 8 | 4,222 | 3,270 | 3,108 | 1,333 | 899 |
| | | 16 | 8,738 | 6,940 | 6,302 | 2,526 | 1,806 |
| | RAID5 | 3 | 2,056 | 1,196 | 1,106 | 454 | 448 |
| | | 4 | 2,327 | 1,476 | 1,316 | 663 | 672 |
| | | 8 | 3,220 | 2,334 | 1,962 | 1,573 | 1,562 |
| | | 16 | 6,919 | 5,056 | 4,344 | 3,316 | 3,336 |
| PRAID EP540i | Single | 1 | 607 | 471 | 463 | 232 | 229 |
| WT: | RAID0 | 2 | 1,296 | 975 | 958 | 550 | 455 |
| Write Through | | 4 | 2,607 | 1,903 | 1,965 | 960 | 894 |
| | | 8 | 5,336 | 3,922 | 3,980 | 1,831 | 1,474 |
| | | 16 | 11,403 | 8,519 | 8,221 | 3,592 | 3,280 |
| | RAID1/10 | 2 | 939 | 701 | 597 | 250 | 228 |
| | | 4 | 1,958 | 1,480 | 1,278 | 510 | 445 |
| | | 8 | 4,116 | 3,032 | 2,607 | 1,017 | 887 |
| | | 16 | 8,489 | 6,500 | 5,537 | 1,900 | 1,731 |
| | RAID5 | 3 | 1,131 | 849 | 713 | 454 | 447 |
| | | 4 | 1,291 | 986 | 776 | 664 | 668 |
| | | 8 | 2,600 | 1,988 | 1,532 | 1,616 | 1,410 |
| | | 16 | 5,588 | 4,111 | 3,332 | 3,422 | 1,641 |
| PRAID EP540i | Single | 1 | 687 | 600 | 809 | 232 | 229 |
| AWB: | RAID0 | 2 | 1,309 | 1,033 | 1,321 | 461 | 470 |
| Always Write Back | | 4 | 3,386 | 2,295 | 2,372 | 885 | 897 |
| | | 8 | 6,089 | 4,438 | 4,600 | 1,837 | 1,788 |
| | | 16 | 11,392 | 9,208 | 9,889 | 3,552 | 3,561 |
| | RAID1/10 | 2 | 1,059 | 987 | 1,220 | 359 | 212 |
| | | 4 | 2,021 | 1,691 | 1,644 | 693 | 453 |
| | | 8 | 4,331 | 3,443 | 3,214 | 1,331 | 900 |
| | | 16 | 9,021 | 7,074 | 6,775 | 2,511 | 1,808 |
| | RAID5 | 3 | 2,102 | 1,752 | 1,328 | 455 | 446 |
| | | 4 | 2,314 | 1,947 | 1,519 | 667 | 671 |
| | | 8 | 4,423 | 2,612 | 2,180 | 1,615 | 1,563 |
| | | 16 | 7,442 | 5,671 | 5,040 | 3,398 | 3,334 |
| PRAID EP640i | Single | 1 | 608 | 478 | 461 | 232 | 228 |
| WT: | RAID0 | 2 | 1,289 | 986 | 956 | 552 | 452 |
| Write Through | | 4 | 2,659 | 1,985 | 1,974 | 964 | 884 |
| | | 8 | 5,433 | 3,954 | 3,903 | 1,845 | 1,744 |
| | | 16 | 11,348 | 8,130 | 7,952 | 3,601 | 3,247 |
| | RAID1/10 | 2 | 945 | 683 | 590 | 251 | 227 |
| | | 4 | 1,906 | 1,442 | 1,257 | 502 | 446 |
| | | 8 | 3,995 | 2,933 | 2,560 | 1,019 | 779 |
| | | 16 | 8,495 | 6,147 | 5,452 | 1,893 | 1,128 |
| | RAID5 | 3 | 1,119 | 852 | 712 | 425 | 439 |
| | | 4 | 1,264 | 976 | 755 | 666 | 665 |
| | | 8 | 2,582 | 1,912 | 1,496 | 1,625 | 1,395 |
| | | 16 | 5,606 | 4,104 | 3,247 | 3,422 | 1,692 |
| PRAID EP640i | Single | 1 | 721 | 598 | 808 | 232 | 229 |
| AWB: | RAID0 | 2 | 1,280 | 1,020 | 1,293 | 462 | 457 |
| Always Write Back | | 4 | 3,380 | 2,290 | 2,366 | 884 | 896 |
| | | 8 | 5,763 | 4,318 | 4,656 | 1,852 | 1,790 |
| | | 16 | 11,896 | 9,221 | 9,675 | 3,654 | 3,562 |
| | RAID1/10 | 2 | 1,097 | 980 | 1,310 | 376 | 228 |
| | | 4 | 2,020 | 1,728 | 1,635 | 703 | 453 |
| | | 8 | 4,298 | 3,377 | 3,183 | 1,333 | 900 |
| | | 16 | 8,924 | 7,068 | 6,699 | 2,527 | 1,812 |
| | RAID5 | 3 | 2,080 | 1,683 | 1,439 | 454 | 447 |
| | | 4 | 2,295 | 1,911 | 1,508 | 664 | 670 |
| | | 8 | 4,181 | 2,572 | 2,047 | 1,547 | 1,563 |
| | | 16 | 7,686 | 5,541 | 4,821 | 3,437 | 3,342 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID EP680i WT: Write Through | Single | 1 | 604 | 475 | 467 | 232 | 229 |
| | RAID0 | 2 | 1,328 | 1,009 | 998 | 652 | 457 |
| | | 4 | 2,789 | 2,077 | 2,048 | 1,062 | 895 |
| | | 8 | 5,672 | 4,079 | 4,068 | 1,927 | 1,750 |
| | | 16 | 11,364 | 8,263 | 8,193 | 3,723 | 3,223 |
| | RAID1/10 | 2 | 936 | 686 | 589 | 252 | 227 |
| | | 4 | 1,907 | 1,428 | 1,252 | 500 | 444 |
| | | 8 | 4,139 | 2,951 | 2,504 | 1,113 | 721 |
| | | 16 | 8,393 | 6,339 | 5,602 | 1,819 | 1,720 |
| | RAID5 | 3 | 1,132 | 863 | 707 | 454 | 448 |
| | | 4 | 1,296 | 974 | 788 | 665 | 669 |
| | | 8 | 2,586 | 1,964 | 1,544 | 1,719 | 1,401 |
| | | 16 | 5,637 | 4,107 | 3,222 | 3,404 | 1,673 |
| PRAID EP680i AWB: Always Write Back | Single | 1 | 756 | 632 | 850 | 232 | 229 |
| | RAID0 | 2 | 1,337 | 1,088 | 1,349 | 461 | 459 |
| | | 4 | 3,923 | 2,485 | 3,233 | 882 | 897 |
| | | 8 | 7,120 | 4,960 | 5,326 | 1,931 | 1,818 |
| | | 16 | 11,063 | 9,934 | 11,156 | 3,779 | 3,570 |
| | RAID1/10 | 2 | 1,133 | 994 | 1,343 | 377 | 229 |
| | | 4 | 2,076 | 1,738 | 1,670 | 699 | 448 |
| | | 8 | 4,326 | 3,401 | 3,231 | 1,331 | 900 |
| | | 16 | 8,923 | 7,322 | 6,973 | 2,521 | 1,885 |
| | RAID5 | 3 | 2,063 | 1,750 | 1,520 | 454 | 448 |
| | | 4 | 2,330 | 2,004 | 1,495 | 663 | 672 |
| | | 8 | 6,116 | 3,868 | 2,751 | 1,730 | 1,561 |
| | | 16 | 9,483 | 6,292 | 5,686 | 3,579 | 3,345 |
| PRAID EP3254-8i WT: Write Through | Single | 1 | 627 | 471 | 463 | 231 | 230 |
| | RAID0 | 2 | 1,293 | 970 | 955 | 460 | 457 |
| | | 4 | 2,674 | 1,954 | 1,915 | 872 | 894 |
| | | 8 | 5,480 | 3,914 | 4,024 | 1,742 | 1,765 |
| | | 16 | 11,444 | 8,564 | 8,506 | 3,461 | 3,488 |
| | RAID1/10 | 2 | 933 | 689 | 600 | 341 | 228 |
| | | 4 | 1,897 | 1,464 | 1,254 | 684 | 446 |
| | | 8 | 3,977 | 3,025 | 2,542 | 1,211 | 863 |
| | | 16 | 8,536 | 6,432 | 5,457 | 2,011 | 809 |
| | RAID5 | 3 | 1,149 | 881 | 710 | 568 | 447 |
| | | 4 | 1,306 | 999 | 763 | 791 | 598 |
| | | 8 | 2,658 | 2,017 | 1,585 | 1,672 | 459 |
| | | 16 | 5,607 | 4,236 | 3,294 | 3,347 | 565 |
| PRAID EP3254-8i AWB: Always Write Back | Single | 1 | 893 | 731 | 603 | 230 | 230 |
| | RAID0 | 2 | 1,896 | 1,493 | 1,944 | 462 | 459 |
| | | 4 | 4,019 | 2,872 | 3,048 | 873 | 898 |
| | | 8 | 6,490 | 5,463 | 4,628 | 1,748 | 1,782 |
| | | 16 | 11,350 | 9,550 | 9,678 | 3,472 | 3,503 |
| | RAID1/10 | 2 | 1,849 | 1,544 | 855 | 408 | 229 |
| | | 4 | 3,977 | 3,258 | 2,962 | 796 | 449 |
| | | 8 | 4,212 | 4,033 | 4,507 | 1,389 | 895 |
| | | 16 | 10,420 | 7,367 | 7,340 | 2,201 | 1,785 |
| | RAID5 | 3 | 2,698 | 1,095 | 1,216 | 569 | 446 |
| | | 4 | 2,991 | 907 | 1,480 | 791 | 670 |
| | | 8 | 4,098 | 1,993 | 2,299 | 1,672 | 1,561 |
| | | 16 | 6,094 | 3,702 | 3,303 | 3,370 | 3,275 |
| PRAID EP3258-16i WT: Write Through | Single | 1 | 616 | 466 | 457 | 231 | 230 |
| | RAID0 | 2 | 1,241 | 942 | 916 | 460 | 455 |
| | | 4 | 2,550 | 1,891 | 1,906 | 879 | 888 |
| | | 8 | 5,210 | 3,842 | 3,905 | 1,751 | 1,608 |
| | | 16 | 11,162 | 8,287 | 8,417 | 3,448 | 3,452 |
| | RAID1/10 | 2 | 917 | 697 | 609 | 410 | 229 |
| | | 4 | 1,929 | 1,479 | 1,250 | 794 | 447 |
| | | 8 | 3,989 | 2,990 | 2,570 | 1,397 | 868 |
| | | 16 | 8,547 | 6,326 | 5,406 | 2,165 | 790 |
| | RAID5 | 3 | 1,122 | 846 | 709 | 570 | 449 |
| | | 4 | 1,256 | 947 | 736 | 788 | 624 |
| | | 8 | 2,548 | 1,961 | 1,544 | 1,673 | 490 |
| | | 16 | 5,468 | 4,080 | 3,178 | 3,377 | 562 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID EP3258-16i | Single | 1 | 875 | 702 | 613 | 227 | 229 |
| AWB: | RAID0 | 2 | 1,834 | 1,384 | 1,594 | 458 | 458 |
| Always Write Back | | 4 | 3,842 | 2,921 | 3,634 | 876 | 893 |
| | | 8 | 6,880 | 5,916 | 5,740 | 1,750 | 1,757 |
| | | 16 | 11,343 | 11,122 | 11,289 | 3,471 | 3,482 |
| | RAID1/10 | 2 | 1,808 | 1,493 | 914 | 410 | 229 |
| | | 4 | 3,920 | 3,188 | 2,059 | 806 | 446 |
| | | 8 | 4,200 | 4,331 | 4,123 | 1,385 | 891 |
| | | 16 | 10,944 | 7,790 | 8,089 | 2,068 | 1,739 |
| | RAID5 | 3 | 2,518 | 1,581 | 1,185 | 571 | 448 |
| | | 4 | 3,338 | 1,387 | 1,810 | 790 | 667 |
| | | 8 | 6,839 | 2,108 | 2,967 | 1,671 | 1,543 |
| | | 16 | 7,599 | 4,346 | 3,951 | 3,339 | 3,193 |

 

## SSD connection

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PSAS CP2100-8i | Single | 1 | 106,142 | 20,548 | 18,151 | 1,032 | 843 |
| | RAID0 | 2 | 211,739 | 40,103 | 35,944 | 2,065 | 1,660 |
| | | 4 | 247,598 | 78,654 | 69,252 | 4,134 | 3,328 |
| | | 8 | 248,123 | 140,794 | 131,356 | 6,805 | 6,310 |
| | | 16 | 267,896 | 146,798 | 140,439 | 6,779 | 6,841 |
| | RAID1/10 | 2 | 153,015 | 27,211 | 21,578 | 1,322 | 831 |
| | | 4 | 262,362 | 53,477 | 42,557 | 2,646 | 1,608 |
| | | 8 | 265,340 | 100,269 | 80,926 | 5,311 | 3,193 |
| | | 16 | 244,353 | 108,284 | 89,112 | 6,809 | 3,424 |
| | RAID5 | 3 | 75,398 | 9,670 | 6,398 | 2,508 | 348 |
| | | 4 | 58,846 | 7,557 | 4,991 | 3,582 | 377 |
| | | 8 | 58,883 | 7,543 | 4,978 | 6,780 | 419 |
| | | 16 | 58,847 | 7,532 | 4,977 | 6,802 | 439 |
| PSAS CP2200-16i | Single | 1 | 133,361 | 23,955 | 19,628 | 1,045 | 882 |
| | RAID0 | 2 | 231,689 | 47,440 | 39,002 | 2,087 | 1,767 |
| | | 4 | 260,983 | 94,284 | 77,913 | 4,177 | 3,542 |
| | | 8 | 267,609 | 176,488 | 140,193 | 8,335 | 6,893 |
| | | 16 | 265,864 | 200,781 | 201,535 | 12,280 | 11,214 |
| | RAID1/10 | 2 | 196,759 | 29,481 | 22,222 | 1,436 | 886 |
| | | 4 | 258,626 | 58,832 | 45,538 | 2,866 | 1,768 |
| | | 8 | 263,334 | 116,688 | 90,479 | 5,566 | 3,527 |
| | | 16 | 265,154 | 190,628 | 144,365 | 10,135 | 5,856 |
| | RAID5 | 3 | 111,232 | 13,968 | 9,271 | 2,507 | 658 |
| | | 4 | 77,511 | 9,780 | 6,472 | 3,549 | 722 |
| | | 8 | 77,349 | 9,774 | 6,433 | 7,655 | 808 |
| | | 16 | 77,419 | 9,732 | 6,444 | 11,211 | 849 |
| PRAID CP400i | Single | 1 | 122,926 | 20,045 | 17,826 | 1,050 | 981 |
| | RAID0 | 2 | 141,107 | 39,201 | 35,291 | 2,007 | 1,700 |
| | | 4 | 146,846 | 74,907 | 68,095 | 3,864 | 3,371 |
| | | 8 | 148,405 | 117,768 | 119,148 | 5,866 | 5,135 |
| | | 16 | 147,216 | 115,152 | 126,130 | 5,905 | 6,469 |
| | RAID1/10 | 2 | 102,454 | 26,806 | 21,752 | 1,410 | 977 |
| | | 4 | 99,145 | 52,048 | 42,597 | 3,360 | 1,695 |
| | | 8 | 94,574 | 89,859 | 74,725 | 4,870 | 2,669 |
| | | 16 | 97,379 | 93,090 | 82,093 | 5,915 | 3,262 |
| | RAID5 | 3 | 24,628 | 23,521 | 18,202 | 3,025 | 1,552 |
| | | 4 | 21,804 | 21,598 | 15,289 | 3,857 | 1,890 |
| | | 8 | 22,374 | 22,158 | 15,665 | 5,874 | 1,975 |
| | | 16 | 22,270 | 21,409 | 15,816 | 5,903 | 1,994 |
| PRAID CP500i | Single | 1 | 122,829 | 22,915 | 19,803 | 1,044 | 856 |
| | RAID0 | 2 | 243,848 | 45,270 | 39,244 | 2,085 | 1,710 |
| | | 4 | 297,721 | 89,023 | 78,157 | 4,171 | 3,412 |
| | | 8 | 282,813 | 147,029 | 139,454 | 7,052 | 6,564 |
| | | 16 | 305,327 | 155,304 | 142,847 | 7,070 | 6,970 |
| | RAID1/10 | 2 | 178,125 | 29,698 | 22,771 | 1,865 | 857 |
| | | 4 | 267,659 | 59,287 | 45,410 | 3,348 | 1,652 |
| | | 8 | 279,259 | 106,458 | 85,678 | 6,615 | 2,311 |
| | | 16 | 276,908 | 117,672 | 95,724 | 7,063 | 2,317 |
| | RAID5 | 3 | 82,048 | 38,952 | 29,266 | 3,108 | 1,650 |
| | | 4 | 70,677 | 45,638 | 33,691 | 4,141 | 2,046 |
| | | 8 | 73,855 | 51,184 | 34,745 | 7,025 | 2,184 |
| | | 16 | 73,622 | 50,808 | 34,867 | 7,061 | 2,236 |
| PRAID CP600i | Single | 1 | 126,226 | 20,221 | 18,044 | 1,044 | 875 |
| | RAID0 | 2 | 252,036 | 40,196 | 36,023 | 2,084 | 1,753 |
| | | 4 | 309,911 | 79,705 | 71,487 | 4,171 | 3,496 |
| | | 8 | 287,801 | 114,369 | 123,188 | 8,044 | 6,867 |
| | | 16 | 293,834 | 134,987 | 142,628 | 8,326 | 8,090 |
| | RAID1/10 | 2 | 185,902 | 27,173 | 21,872 | 1,863 | 876 |
| | | 4 | 277,998 | 54,053 | 43,734 | 3,500 | 1,746 |
| | | 8 | 283,883 | 106,428 | 84,462 | 6,791 | 3,480 |
| | | 16 | 276,074 | 112,505 | 98,204 | 8,307 | 4,039 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID EP400i | Single | 1 | 128,887 | 20,048 | 17,865 | 1,041 | 857 |
| WT: | RAID0 | 2 | 250,919 | 39,173 | 35,488 | 2,079 | 1,720 |
| Write Through | | 4 | 286,995 | 77,563 | 69,710 | 4,143 | 3,337 |
| | | 8 | 305,907 | 120,381 | 122,936 | 5,864 | 5,669 |
| | | 16 | 275,701 | 121,748 | 128,378 | 5,856 | 6,557 |
| | RAID1/10 | 2 | 147,207 | 26,695 | 21,618 | 1,498 | 968 |
| | | 4 | 270,800 | 52,567 | 42,263 | 3,511 | 1,670 |
| | | 8 | 303,606 | 93,048 | 73,661 | 5,847 | 2,705 |
| | | 16 | 274,420 | 95,426 | 86,430 | 5,858 | 3,247 |
| | RAID5 | 3 | 126,155 | 34,004 | 26,588 | 3,098 | 1,734 |
| | | 4 | 44,511 | 38,913 | 29,259 | 3,702 | 2,380 |
| | | 8 | 150,545 | 52,085 | 34,373 | 5,864 | 2,941 |
| | | 16 | 151,351 | 52,329 | 34,699 | 5,864 | 3,120 |
| PRAID EP400i | Single | 1 | 68,555 | 20,305 | 18,079 | 1,045 | 975 |
| AWB: | RAID0 | 2 | 72,416 | 39,994 | 35,999 | 2,002 | 1,711 |
| Always Write Back | | 4 | 72,564 | 73,769 | 61,422 | 3,848 | 3,358 |
| | | 8 | 72,376 | 74,535 | 61,655 | 5,853 | 5,363 |
| | | 16 | 71,750 | 74,465 | 61,244 | 5,864 | 6,482 |
| | RAID1/10 | 2 | 60,127 | 27,355 | 21,940 | 1,642 | 976 |
| | | 4 | 59,145 | 54,066 | 43,915 | 3,362 | 1,724 |
| | | 8 | 58,856 | 60,275 | 48,307 | 5,822 | 2,652 |
| | | 16 | 57,542 | 59,423 | 47,559 | 5,891 | 3,251 |
| | RAID5 | 3 | 41,526 | 34,316 | 26,869 | 3,016 | 1,714 |
| | | 4 | 37,300 | 35,773 | 26,625 | 3,852 | 2,329 |
| | | 8 | 38,460 | 37,415 | 27,323 | 5,877 | 2,831 |
| | | 16 | 40,426 | 37,673 | 27,569 | 5,888 | 3,007 |
| PRAID EP420i | Single | 1 | 128,938 | 20,063 | 17,757 | 1,041 | 863 |
| WT: | RAID0 | 2 | 249,835 | 39,275 | 35,387 | 2,077 | 1,724 |
| Write Through | | 4 | 276,314 | 77,411 | 69,826 | 4,146 | 3,433 |
| | | 8 | 275,445 | 120,141 | 121,473 | 5,860 | 5,651 |
| | | 16 | 304,281 | 121,701 | 128,507 | 5,866 | 6,554 |
| | RAID1/10 | 2 | 185,714 | 26,582 | 21,595 | 1,739 | 969 |
| | | 4 | 293,527 | 52,066 | 42,936 | 3,462 | 1,610 |
| | | 8 | 274,276 | 90,914 | 72,259 | 5,847 | 2,612 |
| | | 16 | 271,568 | 95,865 | 86,759 | 5,857 | 3,250 |
| | RAID5 | 3 | 124,038 | 33,912 | 26,468 | 3,098 | 1,730 |
| | | 4 | 104,781 | 38,727 | 28,779 | 4,147 | 2,279 |
| | | 8 | 149,668 | 51,124 | 33,831 | 5,860 | 2,903 |
| | | 16 | 151,285 | 51,238 | 33,987 | 5,862 | 3,056 |
| PRAID EP420i | Single | 1 | 66,876 | 20,528 | 18,237 | 1,047 | 977 |
| AWB: | RAID0 | 2 | 70,448 | 40,425 | 36,227 | 2,000 | 1,704 |
| Always Write Back | | 4 | 70,684 | 73,264 | 61,189 | 3,862 | 3,348 |
| | | 8 | 70,578 | 74,485 | 61,471 | 5,845 | 5,305 |
| | | 16 | 69,532 | 73,700 | 61,187 | 5,891 | 6,430 |
| | RAID1/10 | 2 | 58,615 | 28,059 | 21,806 | 1,654 | 918 |
| | | 4 | 57,954 | 54,486 | 43,712 | 3,422 | 1,714 |
| | | 8 | 40,514 | 53,580 | 42,598 | 5,194 | 2,722 |
| | | 16 | 56,140 | 58,932 | 47,262 | 5,873 | 3,243 |
| | RAID5 | 3 | 40,901 | 34,760 | 27,221 | 3,016 | 1,709 |
| | | 4 | 37,255 | 36,133 | 26,542 | 3,885 | 2,197 |
| | | 8 | 38,870 | 37,433 | 27,473 | 5,840 | 2,780 |
| | | 16 | 42,277 | 38,331 | 28,149 | 5,887 | 2,954 |
| PRAID EP520i | Single | 1 | 122,479 | 22,844 | 19,830 | 1,043 | 857 |
| WT: | RAID0 | 2 | 242,888 | 45,347 | 39,218 | 2,085 | 1,710 |
| Write Through | | 4 | 279,925 | 89,462 | 77,757 | 4,165 | 3,406 |
| | | 8 | 287,457 | 146,521 | 139,091 | 7,048 | 6,511 |
| | | 16 | 291,734 | 153,556 | 153,804 | 7,063 | 6,929 |
| | RAID1/10 | 2 | 177,081 | 29,706 | 22,707 | 1,862 | 866 |
| | | 4 | 276,097 | 59,090 | 45,436 | 3,398 | 1,720 |
| | | 8 | 267,193 | 105,661 | 85,124 | 6,661 | 3,381 |
| | | 16 | 265,025 | 116,117 | 94,633 | 7,076 | 3,878 |
| | RAID5 | 3 | 73,409 | 39,151 | 29,588 | 3,110 | 1,733 |
| | | 4 | 63,344 | 45,301 | 34,641 | 4,147 | 2,583 |
| | | 8 | 87,110 | 62,921 | 41,901 | 7,028 | 3,580 |
| | | 16 | 86,374 | 63,142 | 41,933 | 7,063 | 3,713 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID EP520i | Single | 1 | 116,068 | 23,437 | 20,734 | 1,043 | 864 |
| AWB: | RAID0 | 2 | 120,351 | 44,756 | 40,234 | 2,086 | 1,759 |
| Always Write Back | | 4 | 119,869 | 91,388 | 80,032 | 4,165 | 3,465 |
| | | 8 | 120,364 | 123,842 | 88,670 | 7,061 | 3,688 |
| | | 16 | 121,613 | 123,408 | 88,630 | 7,060 | 3,747 |
| | RAID1/10 | 2 | 90,970 | 30,311 | 23,262 | 1,866 | 858 |
| | | 4 | 89,387 | 61,090 | 47,031 | 3,411 | 1,715 |
| | | 8 | 88,601 | 99,914 | 70,360 | 6,661 | 3,265 |
| | | 16 | 87,930 | 99,375 | 69,677 | 7,063 | 2,371 |
| | RAID5 | 3 | 75,959 | 39,489 | 29,782 | 2,623 | 1,607 |
| | | 4 | 64,264 | 46,051 | 34,980 | 4,152 | 2,584 |
| | | 8 | 62,546 | 58,788 | 40,237 | 7,040 | 3,614 |
| | | 16 | 67,064 | 57,071 | 38,750 | 7,063 | 3,713 |
| PRAID EP540i | Single | 1 | 134,494 | 23,709 | 19,479 | 1,052 | 877 |
| WT: | RAID0 | 2 | 268,245 | 47,232 | 38,958 | 2,103 | 1,753 |
| Write Through | | 4 | 284,069 | 84,461 | 70,883 | 4,199 | 3,415 |
| | | 8 | 289,854 | 143,761 | 131,897 | 7,050 | 6,584 |
| | | 16 | 304,910 | 164,775 | 194,031 | 7,036 | 6,942 |
| | RAID1/10 | 2 | 198,453 | 29,277 | 22,705 | 2,102 | 887 |
| | | 4 | 266,660 | 57,369 | 45,288 | 3,369 | 1,734 |
| | | 8 | 278,278 | 109,254 | 85,270 | 6,664 | 2,849 |
| | | 16 | 265,592 | 144,441 | 108,336 | 7,065 | 4,094 |
| | RAID5 | 3 | 73,263 | 39,579 | 29,362 | 3,154 | 1,763 |
| | | 4 | 63,098 | 46,777 | 34,127 | 4,200 | 2,636 |
| | | 8 | 86,094 | 62,127 | 41,344 | 7,028 | 3,509 |
| | | 16 | 85,555 | 62,213 | 41,178 | 7,066 | 3,540 |
| PRAID EP540i | Single | 1 | 107,754 | 25,171 | 21,897 | 1,052 | 882 |
| AWB: | RAID0 | 2 | 109,017 | 47,731 | 41,259 | 2,102 | 1,838 |
| Always Write Back | | 4 | 108,943 | 97,049 | 78,863 | 4,200 | 3,518 |
| | | 8 | 110,146 | 110,417 | 83,353 | 7,050 | 3,691 |
| | | 16 | 112,722 | 110,350 | 83,403 | 7,064 | 3,729 |
| | RAID1/10 | 2 | 84,789 | 30,852 | 23,928 | 2,099 | 885 |
| | | 4 | 83,406 | 61,585 | 47,726 | 3,444 | 1,747 |
| | | 8 | 82,738 | 91,427 | 67,212 | 6,728 | 2,953 |
| | | 16 | 82,002 | 90,545 | 66,327 | 7,066 | 2,361 |
| | RAID5 | 3 | 73,112 | 40,685 | 30,516 | 3,155 | 1,773 |
| | | 4 | 63,295 | 48,823 | 36,066 | 4,204 | 2,660 |
| | | 8 | 61,210 | 59,535 | 40,427 | 7,038 | 3,592 |
| | | 16 | 71,524 | 59,854 | 40,608 | 7,065 | 3,542 |
| PRAID EP640i | Single | 1 | 126,454 | 20,193 | 18,039 | 1,043 | 873 |
| WT: | RAID0 | 2 | 251,719 | 40,177 | 35,941 | 2,087 | 1,743 |
| Write Through | | 4 | 275,000 | 79,554 | 71,581 | 4,171 | 3,474 |
| | | 8 | 284,802 | 112,878 | 122,324 | 8,226 | 6,689 |
| | | 16 | 280,034 | 147,462 | 143,780 | 8,325 | 8,033 |
| | RAID1/10 | 2 | 183,208 | 27,152 | 21,903 | 1,861 | 818 |
| | | 4 | 276,974 | 53,796 | 43,738 | 3,416 | 1,731 |
| | | 8 | 265,509 | 103,510 | 85,362 | 6,635 | 3,223 |
| | | 16 | 270,683 | 121,784 | 98,415 | 8,307 | 4,041 |
| | RAID5 | 3 | 187,887 | 34,933 | 27,258 | 3,100 | 1,757 |
| | | 4 | 250,140 | 40,539 | 30,747 | 4,158 | 2,609 |
| | | 8 | 274,137 | 77,834 | 57,160 | 8,306 | 4,832 |
| | | 16 | 274,909 | 85,812 | 57,429 | 8,324 | 5,139 |
| PRAID EP640i | Single | 1 | 128,956 | 21,883 | 18,998 | 1,043 | 887 |
| AWB: | RAID0 | 2 | 146,594 | 43,426 | 37,727 | 2,086 | 1,766 |
| Always Write Back | | 4 | 146,154 | 84,336 | 75,510 | 4,168 | 3,426 |
| | | 8 | 146,680 | 119,905 | 113,251 | 8,288 | 4,866 |
| | | 16 | 149,546 | 147,540 | 112,996 | 8,325 | 5,097 |
| | RAID1/10 | 2 | 112,465 | 28,507 | 23,049 | 1,864 | 887 |
| | | 4 | 110,816 | 57,858 | 46,147 | 3,510 | 1,763 |
| | | 8 | 109,397 | 98,147 | 71,338 | 6,643 | 2,518 |
| | | 16 | 108,671 | 97,948 | 70,944 | 8,270 | 2,485 |
| | RAID5 | 3 | 95,084 | 36,493 | 27,081 | 3,115 | 1,770 |
| | | 4 | 82,382 | 43,461 | 32,781 | 3,986 | 2,656 |
| | | 8 | 79,824 | 74,987 | 52,457 | 8,215 | 4,722 |
| | | 16 | 91,683 | 75,018 | 53,227 | 8,325 | 5,205 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID EP680i WT: Write Through | Single | 1 | 136,415 | 23,791 | 19,518 | 1,052 | 880 |
| | RAID0 | 2 | 269,477 | 47,227 | 39,093 | 2,104 | 1,757 |
| | | 4 | 284,489 | 92,260 | 77,680 | 4,206 | 3,502 |
| | | 8 | 291,944 | 176,044 | 152,008 | 8,385 | 6,988 |
| | | 16 | 288,110 | 232,800 | 223,623 | 13,762 | 13,334 |
| | RAID1/10 | 2 | 198,253 | 29,550 | 22,816 | 2,101 | 876 |
| | | 4 | 287,336 | 58,652 | 45,427 | 3,371 | 1,755 |
| | | 8 | 275,871 | 114,832 | 89,192 | 6,640 | 3,506 |
| | | 16 | 277,427 | 205,177 | 165,809 | 13,242 | 6,650 |
| | RAID5 | 3 | 201,994 | 39,663 | 29,414 | 3,152 | 1,780 |
| | | 4 | 259,678 | 47,275 | 34,780 | 4,200 | 2,643 |
| | | 8 | 280,608 | 83,212 | 56,026 | 8,397 | 4,810 |
| | | 16 | 287,720 | 85,413 | 56,805 | 14,096 | 4,983 |
| PRAID EP680i AWB: Always Write Back | Single | 1 | 134,437 | 26,302 | 21,575 | 1,052 | 893 |
| | RAID0 | 2 | 138,567 | 51,679 | 43,666 | 2,104 | 1,783 |
| | | 4 | 138,427 | 100,401 | 85,663 | 4,201 | 3,560 |
| | | 8 | 139,324 | 130,861 | 105,662 | 8,346 | 5,033 |
| | | 16 | 141,154 | 131,671 | 106,637 | 13,699 | 5,128 |
| | RAID1/10 | 2 | 108,019 | 32,493 | 25,187 | 2,103 | 895 |
| | | 4 | 106,812 | 64,852 | 50,247 | 3,445 | 1,785 |
| | | 8 | 105,995 | 90,946 | 68,572 | 6,831 | 2,509 |
| | | 16 | 104,840 | 91,017 | 68,349 | 13,371 | 2,495 |
| | RAID5 | 3 | 94,238 | 45,212 | 32,805 | 3,151 | 1,706 |
| | | 4 | 84,906 | 54,736 | 40,419 | 4,201 | 2,676 |
| | | 8 | 82,277 | 76,510 | 55,830 | 8,397 | 4,890 |
| | | 16 | 91,832 | 82,091 | 57,758 | 13,744 | 4,923 |
| PRAID EP3254-8i WT: Write Through | Single | 1 | 131,278 | 23,735 | 19,421 | 1,040 | 876 |
| | RAID0 | 2 | 251,426 | 46,509 | 38,133 | 2,080 | 1,751 |
| | | 4 | 242,212 | 91,258 | 72,339 | 4,152 | 3,488 |
| | | 8 | 267,138 | 181,449 | 131,641 | 8,267 | 6,719 |
| | | 16 | 266,166 | 185,270 | 159,787 | 8,281 | 7,271 |
| | RAID1/10 | 2 | 195,861 | 28,980 | 22,454 | 1,424 | 879 |
| | | 4 | 262,929 | 55,268 | 42,880 | 2,864 | 1,751 |
| | | 8 | 241,880 | 96,504 | 74,609 | 5,492 | 3,462 |
| | | 16 | 241,652 | 116,309 | 85,756 | 8,309 | 3,605 |
| | RAID5 | 3 | 188,978 | 38,130 | 28,209 | 2,512 | 1,759 |
| | | 4 | 249,237 | 43,817 | 31,692 | 3,591 | 2,516 |
| | | 8 | 266,298 | 62,705 | 41,641 | 7,768 | 4,705 |
| | | 16 | 239,253 | 62,901 | 41,620 | 8,250 | 4,912 |
| PRAID EP3254-8i AWB: Always Write Back | Single | 1 | 97,107 | 24,717 | 20,253 | 1,039 | 967 |
| | RAID0 | 2 | 232,326 | 48,146 | 39,363 | 2,078 | 1,957 |
| | | 4 | 241,298 | 95,062 | 75,369 | 4,155 | 3,946 |
| | | 8 | 244,513 | 138,683 | 100,254 | 8,259 | 6,709 |
| | | 16 | 264,522 | 169,012 | 119,883 | 8,304 | 6,647 |
| | RAID1/10 | 2 | 200,849 | 31,038 | 24,259 | 1,428 | 978 |
| | | 4 | 241,345 | 57,406 | 45,011 | 2,862 | 1,966 |
| | | 8 | 243,185 | 109,508 | 80,113 | 5,347 | 3,781 |
| | | 16 | 261,837 | 127,254 | 96,136 | 8,311 | 3,819 |
| | RAID5 | 3 | 223,332 | 40,954 | 30,620 | 2,503 | 1,755 |
| | | 4 | 224,224 | 47,390 | 33,980 | 3,580 | 2,590 |
| | | 8 | 204,386 | 63,250 | 42,436 | 7,766 | 4,659 |
| | | 16 | 203,713 | 63,911 | 42,852 | 8,306 | 4,769 |
| PRAID EP3258-16i WT: Write Through | Single | 1 | 132,315 | 23,907 | 19,655 | 1,045 | 887 |
| | RAID0 | 2 | 241,202 | 47,518 | 39,083 | 2,089 | 1,763 |
| | | 4 | 245,064 | 94,331 | 77,762 | 4,175 | 3,535 |
| | | 8 | 245,679 | 181,829 | 154,034 | 8,341 | 7,057 |
| | | 16 | 227,158 | 173,638 | 177,075 | 10,576 | 10,959 |
| | RAID1/10 | 2 | 194,359 | 29,538 | 22,860 | 1,415 | 885 |
| | | 4 | 225,944 | 58,809 | 45,630 | 2,868 | 1,769 |
| | | 8 | 227,944 | 116,315 | 90,307 | 5,728 | 3,528 |
| | | 16 | 225,791 | 173,480 | 172,057 | 10,209 | 6,326 |
| | RAID5 | 3 | 197,523 | 39,090 | 29,151 | 2,519 | 1,765 |
| | | 4 | 224,419 | 45,282 | 32,641 | 3,598 | 2,534 |
| | | 8 | 245,448 | 62,773 | 41,741 | 7,811 | 4,778 |
| | | 16 | 226,138 | 63,181 | 41,659 | 10,482 | 5,027 |

 

| RAID controller | RAID level | # drive | Transaction Rate [IOPS] | | | Data Throughput [MiB/s] | |
|---|---|---|---|---|---|---|---|
| | | | Database | Fileserver | Filecopy | Streaming | Restore |
| PRAID EP3258-16i | Single | 1 | 96,205 | 25,687 | 21,660 | 1,045 | 893 |
| AWB: | RAID0 | 2 | 223,482 | 51,019 | 43,914 | 2,088 | 1,780 |
| Always Write Back | | 4 | 225,282 | 102,769 | 84,905 | 4,178 | 3,556 |
| | | 8 | 225,601 | 152,773 | 107,359 | 8,296 | 7,084 |
| | | 16 | 246,694 | 157,717 | 112,519 | 10,663 | 10,073 |
| | RAID1/10 | 2 | 201,259 | 37,253 | 26,334 | 1,445 | 893 |
| | | 4 | 225,001 | 65,268 | 51,498 | 2,850 | 1,780 |
| | | 8 | 225,987 | 127,104 | 97,190 | 5,323 | 3,514 |
| | | 16 | 227,188 | 146,210 | 100,004 | 10,188 | 6,123 |
| | RAID5 | 3 | 214,819 | 46,293 | 34,523 | 2,515 | 1,779 |
| | | 4 | 206,073 | 51,984 | 38,529 | 3,597 | 2,624 |
| | | 8 | 212,835 | 66,439 | 44,805 | 7,805 | 4,755 |
| | | 16 | 214,618 | 68,094 | 45,665 | 11,766 | 4,962 |

 

# Literature

**PRIMERGY & PRIMEQUEST Server**

https://www.fujitsu.com/global/products/computing/servers/

**Performance of Server Components**

https://www.fujitsu.com/global/products/computing/servers/primergy/benchmarks/

**RAID controller performance 2021**

This White Paper:

https://docs.ts.fujitsu.com/dl.aspx?id=454bdca2-90d1-4386-ac28-18fe043e901a

https://docs.ts.fujitsu.com/dl.aspx?id=c5488ccc-aeef-47a4-99c5-3098f2fda66f

**RAID controller performance 2018**

https://docs.ts.fujitsu.com/dl.aspx?id=51032738-c49e-4cd6-ae0f-6de5d79a369d

https://docs.ts.fujitsu.com/dl.aspx?id=db3036e9-12f2-47f1-af66-77d03ca5f773

**Memory performance of systems with Xeon scalable processors**

https://docs.ts.fujitsu.com/dl.aspx?id=1930d389-7521-4c85-bcf9-86a71a14a7c3

**BIOS Optimization for PRIMERGY Xeon Scalable Processor Based RX Systems**

https://docs.ts.fujitsu.com/dl.aspx?id=3cf2d5b2-f58a-4308-ac63-d33adf7230d2

**PRIMERGY BIOS optimizations for Xeon Scalable processors based CX systems**

http://docs.ts.fujitsu.com/dl.aspx?id=e693e228-33b5-48f7-8310-21a313191103

**PRIMEQUEST BIOS optimizations for Xeon Scalable processors based systems**

http://docs.ts.fujitsu.com/dl.aspx?id=4414124c-e3e3-4d04-9e89-fbf8675d1a6b

**RAID Controller Performance 2016 (previous White Paper)**

https://docs.ts.fujitsu.com/dl.aspx?id=9845be50-7d4f-4ef7-ac61-bbde399c1014

**512e HDD: Technology, Performance, Configuration**

https://docs.ts.fujitsu.com/dl.aspx?id=f5550c48-d4db-47f6-ab9d-ce135eaacf81

**Disk I/O Performance Basics**

https://docs.ts.fujitsu.com/dl.aspx?id=65781a00-556f-4a98-90a7-7022feacc602

**Information about Iometer**

https://www.iometer.org

 

## Document change history

| Version | Date | Description |
|---|---|---|
| 1.5 | 2024-07-16 | Update:<br>• Added RAID controller (PSAS CP2200-16i) performance data<br>• Updated RAID controller (PRAID EP3254-8i, PRAID EP3258-16i) performance data<br>• Minor revisions to the description |
| 1.4 | 2023-10-03 | Update:<br>• Support for new VI<br>• Minor corrections |
| 1.3 | 2023-03-27 | Update:<br>• Added RAID controller (PRAID CP600i, PRAID EP640i, PRAID EP3254-8i, PRAID EP3258-16i) performance data<br>• Updated RAID controller (PSAS CP2100-8i) performance data<br>• Minor revisions to the description |
| 1.2 | 2022-09-09 | Update:<br>• Added RAID controller (PRAID CP400i, PRAID EP400i, PRAID EP420i) performance data<br>• Minor revisions to the description |
| 1.1 | 2022-07-05 | Update:<br>• Added RAID controller (VROC NVMe, PSAS CP 2100 -8i) performance data<br>• Updated SAS-SSD performance data of PRAID EP680i<br>• Additional information in "Controller Cache Effects" and "Appendix 2"<br>• Minor revisions to the description |
| 1.0 | 2021-12-02 | New: |

**Contact**
**Fujitsu**
Web site: https://www.fujitsu.com
**PRIMERGY Performance and Benchmarks**
mailto:fj-benchmark@dl.jp.fujitsu.com