Performance Report

PAC Storage 5000U NVMe model
Support 200 GbE

Updated: February 2025

Summary

PAC Storage 5000U is All-Flash U.2 NVMe unified storage system. Equipped with U.2 NVMe SSD, delivering better performance with lower latency.

Contents

1 Applicable Models

For your reference, below are the applicable models in this performance report:

Series	Applicable Models
PS 5024U	PS 5024URE

2 Audience

This performance report is intended for the PAC Storage partners, customers, and employees who want to deploy PAC Storage PS All-Flash U.2 NVMe series as their storage.

3 Terminology

- **PAC Storage PS Family** – PAC Storage PS Family is an enterprise unified storage which can be configured as SAN or NAS.
- **EonOne** – Management software for PAC Storage Systems.
- **Block-level** – Commonly deployed as SAN storage to store structured data.
- **File-level** – Commonly used for storing unstructured data, and shares data with multiple clients via file system protocol, such as CIFS/SMB, NFS and FTP.
- **Logical Drive (LD)** – Combination of multiple drives via RAID technology.
- **Pool** – Combination of one or more LD(s).
- **Volume** – Space divided from pool which can be created as block-level volume or file-level volume.
- **Better performance for block data access** – Assign more CPU cores for block-level IO. Referred as block mode in the rest of this report.
- **Better performance for file access service** – Assign more CPU cores for file-level IO. Referred as file mode in the rest of this report.

4 Performance Results

The following section demonstrates the highest performance of each system. Please also check the configuration in System Configuration section to see how we run the tests. There are also some descriptions of the results in the Conclusion section.

***Color of Performance optimization value:**

Better performance for block data access (Referred as block mode in this report) / Better performance for file access service (Referred as file mode in this report)

4.1 Block-Level Section

4.1.1 IOPS with small block size

Host Type: FC_32G		Block Level				
		Profile	End-to-End			All Cache Hit
		IO Behavior	Random			Sequential
		Size	4KB	8KB	64KB	512B
PS 5024 URE FW: 1.66S.16 Block mode	RAID 5	Read (IOPS)	1,280,777	1,275,345	-	-
		Write (IOPS)	272,777	272,117	-	-
	RAID 6	Read (IOPS)	1,279,820	1,273,612	-	-
		Write (IOPS)	206,222	205,778	-	-

Host Type: iSCSI_25G		Block Level				
		Profile	End-to-End			All Cache Hit
		IO Behavior	Random			Sequential
		Size	4KB	8KB	64KB	512B
PS 5024 URE FW: 1.67F.01 Block mode	RAID 5	Read (IOPS)	818,721	825,647	412,496	1,141,530
		Write (IOPS)	280,767	280,597	97,478	1,064,807

Host Type: iSCSI_100G		Block Level				
		Profile	End-to-End			All Cache Hit
		IO Behavior	Random			Sequential
		Size	4KB	8KB	64KB	512B

PS 5024 URE FW: 1.67F.07 Block mode	RAID 5	Read (IOPS)	1,151,863	1,001,864	422,120	1,259,778
		Write (IOPS)	275,181	274,872	96,033	1,147,105

Host Type: **iSCSI_100G RDMA**		**Block Level**				
		Profile	End-to-End			All Cache Hit
		IO Behavior	Random			Sequential
		Size	4KB	8KB	64KB	512B
PS 5024 URE FW: 1.67F.07 Block mode	RAID 5	Read (IOPS)	1,303,427	1,122,067	720,809	2,306,575
		Write (IOPS)	226,399	230,891	94,187	2,440,460

4.1.2 IOPS with Response Time

Host Type: FC_32G		Block Level					
		Profile		End-to-End			
		Application		Random Read	Random Read	Database R/W = 70%/30%	VDI R/W = 20%/80%
		Response Time	Size	4K	32K	8KB	4KB
PS 5024 URE FW: 1.66S.16 Block mode	RAID 5	< 0.3ms	IOPS	1,280,286	858,411	386,798	178,168
	RAID 6	< 0.3ms	IOPS	1,281,141	859,817	321,278	130,011

Host Type: FC_32G		Block Level					
		Profile		End-to-End			
		Application		Random Read	Random Read	Database R/W = 70%/30%	VDI R/W = 20%/80%
		Response Time	Size	4K	32K	8KB	4KB
PS 5024 URE FW: 1.66S.16 Block mode	RAID 5	< 0.5ms	IOPS	1,280,796	1,039,187	456,736	209,654
	RAID 6	< 0.5ms	IOPS	1,283,695	1,029,954	369,631	151,082

4.1.3 Throughput with large block size

Host Type: SAS_12G		Block Level			
		Profile	End-to-End		All Cache Hit
		IO Behavior	Sequential		Sequential
		Size	64KB	1MB	1MB
PS 5024URE FW: 1.67F.01 Block mode	RAID 5	Read (MB/s)	46,335	46,564	46,800
		Write (MB/s)	16,883	19,446	26,231

Host Type: FC_32G		Block Level			
		Profile	End-to-End		All Cache Hit
		IO Behavior	Sequential		Sequential
		Size	64KB	1MB	1MB
PS 5024URE FW: 166S16 Block mode	RAID 5	Read (MB/s)	-	50,119	50,298
		Write (MB/s)	-	20,047	37,179
	RAID 6	Read (MB/s)	-	50,102	50,298
		Write (MB/s)	-	50,298	37,179

Host Type: iSCSI_25G		Block Level			
		Profile	End-to-End		All Cache Hit
		IO Behavior	Sequential		Sequential
		Size	64KB	1MB	1MB
PS 5024URE FW: 167F01 Block mode	RAID 5	Read (MB/s)	26,403	33,150	33,390
		Write (MB/s)	11,112	15,561	14,615

Host Type: iSCSI_100G RDMA		Block Level			
		Profile	End-to-End		All Cache Hit
		IO Behavior	Sequential		Sequential
		Size	64KB	1MB	1MB
PS 5024URE FW: 1.67F.07 Block mode	RAID 5	Read (MB/s)	46,228	46,483	46,686
		Write (MB/s)	17,579	20,496	38,898

4.1.4 Application Simulation

Host Type: iSCSI_25G		Block Level			
		Profile	End-to-End		
		Application	Database R/W = 70%/30%		VDI R/W = 20%/80%
		Size	4KB	8KB	4KB
PS 5024URE FW: 1.67F.01 Block mode	RAID 5	IOPS	547,505	547,908	313,668

Host Type: iSCSI_100G		Block Level			
		Profile	End-to-End		
		Application	Database R/W = 70%/30%		VDI R/W = 20%/80%
		Size	4KB	8KB	4KB
PS 5024URE FW: 1.67F.07 Block mode	RAID 5	IOPS	550,372	551,080	308,231

4.2 File-Level Section

Host Type: Ethernet_100G RDMA			File Level - NFS (XFS)			
		IO Type	Sequential (MB/s)		Random (IOPS)	
		Size	1MB	512KB	4KB	Database 8KB (R/W: 70%/30%)
PS 5024 URE FW: 1.67F.07 File mode	RAID 5	Read	-	27,016	964,518	235,060
		Write	-	10,502	117,480	

5 Topology

This section illustrated the principle of the network topology and storage configuration. Please refer to topology section and the system configuration section to get the best performance from PAC Storage PS Family. **Note**: To leverage the advantage of multi-thread, please create multiple shared folders to run the file-level tests.

5.1 Block-Level

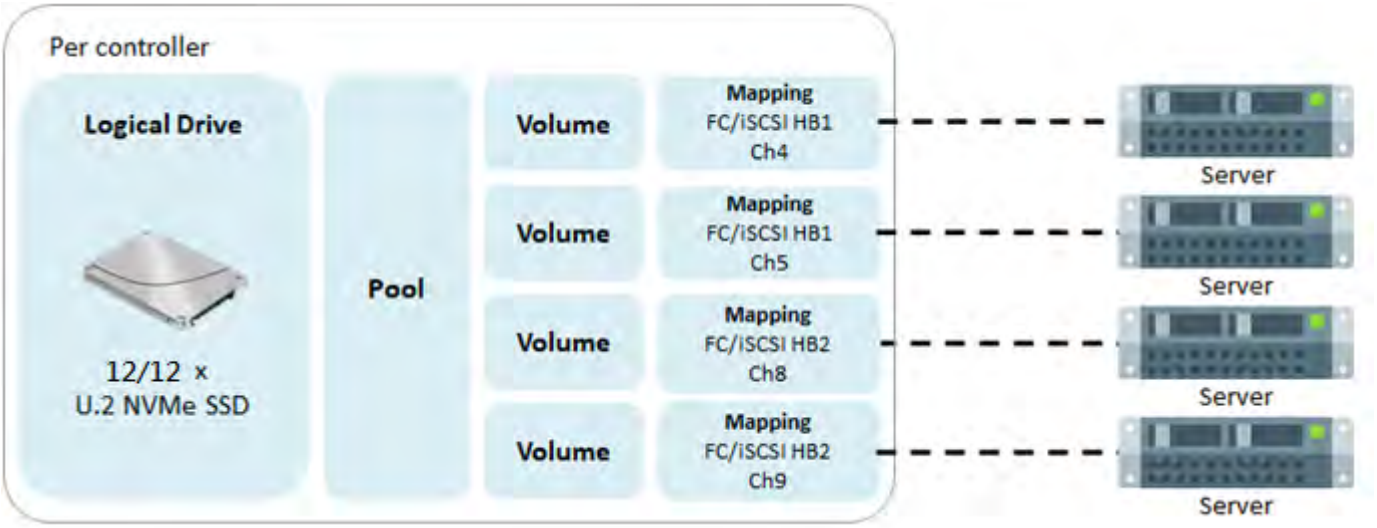

Note: The diagrams above are just for your references. If you need detailed number of channels, please refer to below forms.

5.2 File-Level

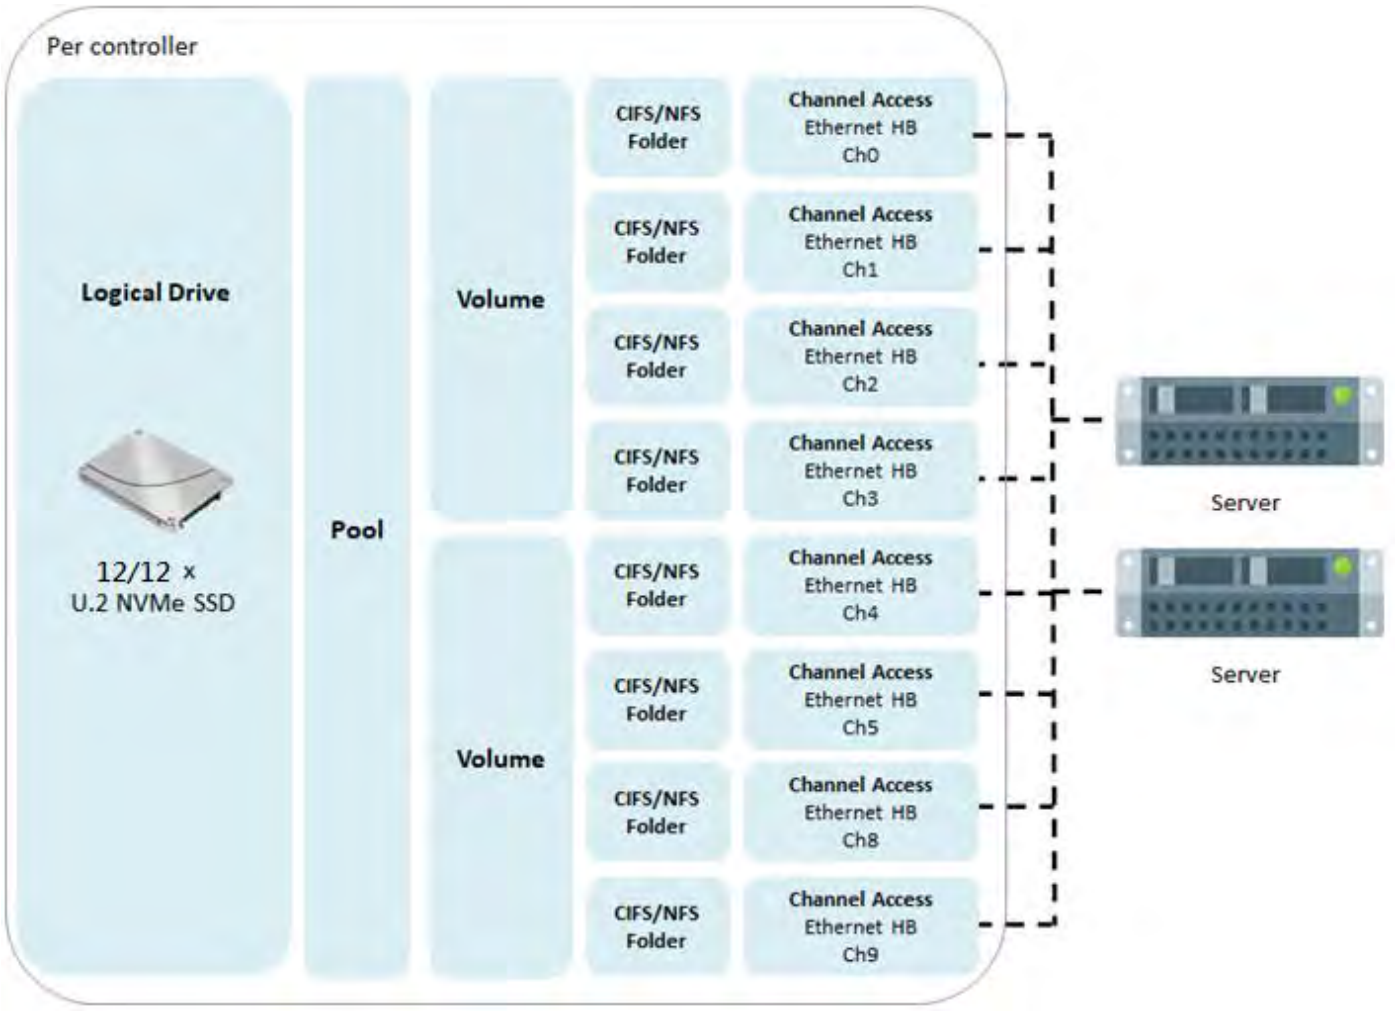

Note: The diagrams above are just for your references. If you need detailed number of channels, please refer to below forms.

Block-level	Model	# of Host Board per controller	# of Channel per controller
SAS 12G	PS 5024URE	3	6
FC 32G		2	8
iSCSI 25G		2	4
iSCSI 100G		2	2
iSCSI 100G RDMA		2	2

File-level	Model	# of Host Board per controller	# of Channel per controller
Ethernet 100G RDMA	PS 5024URE	2	4

6 System Configurations

6.1 Storage Configuration Profile

The following table shows the configuration adopted from our PS best practice with a storage pool and a shared folder. To provide a single namespace sharing solution, we configured the PS dual controller models with an active-standby configuration.

As a tradeoff between usable capacity and failure tolerance, we recommend building the LD within 15 drives.

6.1.1 Block-Level SSD

Model	# of Drive	# of LD	# of Pool	# of Volume	# of Client
PS 5024URE	24	2	2	12	4

6.1.2 File-Level SSD

Model	# of Drive	# of LD	# of Pool	# of Volume	# of Folder	# of Client
PS 5024URE	24	2	2	8	8	4

6.2 Storage System Settings

We use the following parameters to optimize the media workload, which differs from the PS default settings. For detail parameter settings on EonOne, please refer to EonOne software manual.

RAM (per controller)	48GB (16GB x 3)
Stripe size	128K
Maximum Tag Count	64
Periodic Drive Check Time	Disable
Periodic SAF-TE and SES Device Check Time	Disable
Max Drive Response Timeout	Disable
Drive Access Delay Time	No Delay

6.3 Client Workstation Information

The following table shows the specification of the client workstation we used for the performance test. To ensure optimal system performance, we recommend that you deploy a solution with better specifications, especially PCIe lanes and CPU.

M/B	Super Micro X12SPL-F
CPU	Intel Xeon Silver 4309Y CPU 3.60GHZ
RAM	DDR IV 2933 32G*8 (256GB)
PCI	2 PCI-E 4.0 x8, 2 PCI-E 3.0 x8
System Drive	SATA WD 500G (WD5003ABYX-01WERA1)
OS	Windows Server 2019
HBA card	Intel(R) Ethernet Controller E810-C for QSFP (Packet Size:9014) (Receive Buffers:4096) (Transmit Buffers:4096)
MPIO	OS native
Power Option	High Performance
Jumbo Frames	Linux MTU 9000

6.4 Benchmark Tool Settings

Benchmark Tool	**IOmeter 2006.07.27**	
I/O setting	Outstanding	Random – 256, Sequential – 64
	Ramp Up Time	30 sec
	Run Time	180 sec
	All Cache: Maximum Disk Size 102400	
	One LD Corresponds to One Worker.	
	Align I/Os on	

Benchmark Tool	**Vdbench**	
I/O setting	Sequential	Files=10, Size=5g, Threads=10, elapsed=120, interval=10, warmup=20
	Random	Files=10, Size=5g, Threads=64, elapsed=120, interval=10, warmup=20
	All Cache	Files=6, Size=10m, Thread=4, elapsed=120, interval=10, warmup=20
	Random32K(Linux)	Files=125000, Size=32k, Threads=5, elapsed=60, interval=5, warmup=10
	Random32K(Windows)	Files=45000, Size=32k, Threads=5, elapsed=60, interval=5, warmup=10

7 Conclusion

Comparing to SAS/SATA SSD, U.2 NVMe SSD can communicate with CPU via PCIe interface directly, so the latency can be much lower. With lower latency and larger bandwidth, U.2 NVMe SSD can deliver higher performance, and users can reach the performance limit of controller with fewer drives. This makes U.2 NVMe PS a more cost-effective All-Flash solution.

8 Legal Information

PAC Storage may from time to time modify, update or upgrade the software, firmware or any accompanying user documentation without any prior notice. PAC Storage will provide access to these new software, firmware, or documentation releases from certain download sections of our website or through our service partners. Customer will be responsible for maintaining updated version of the software, firmware, or other documentation by downloading or obtaining from PAC Storage, and installing designated updated code, including but not limited to firmware, microcode, basic input/out system code, utility programs, device drivers, and diagnostics delivered with PAC Storage product.
Before installing any software, applications or components provided by a third party, customer should ensure that they are compatible and interoperable with PAC Storage product by checking in advance with PAC Storage. Customer is solely responsible for ensuring the compatibility and interoperability of the third party's products with PAC Storage product.

Customer is further solely responsible for ensuring its systems, software, and data are adequately backed up as a precaution against possible failures, alternation, or loss. For any questions of hardware/ software compatibility, and the update/ upgrade code, customer should contact PAC Storage sales representative or technical support for assistance.

To the extent permitted by applicable laws, PAC Storage will NOT be responsible for any interoperability or compatibility issues that may arise when (1) products, software, or options not certified and supported by PAC Storage are used; (2) configurations not certified and supported by PAC Storage are used; (3) parts intended for one system are installed in another system of different make or model.

9 Contact Information

Website

For more information of PAC Storage's products and services,

visit: https://www.pacstorage.com/