

QSAN NVMe All-Flash Array x Western Digital NVMe SSD Accelerate Applications with Microsecond-Level Latency

XCubeFAS Series Best Practice

May 2021

ANNOUNCEMENT

Copyright

© Copyright 2021 QSAN Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from QSAN Technology, Inc.

QSAN believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Trademarks

- QSAN, the QSAN logo, QSAN.com, XCubeFAS, XCubeSAN, XCubeNXT, XCubeNAS, XCubeDAS, XEVO, SANOS, and QSM are trademarks or registered trademarks of QSAN Technology, Inc.
- Western Digital, the Western Digital logo, ArticFlow, IsoVibe, and Ultrastar are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries.
- Other trademarks and trade names used in this document to refer to either the entities claiming the marks and names or their products are the property of their respective owners.

BEST PRACTICE

Data transmission is fundamental for all kinds of computing. Technologies such as machine learning, edge computing, and virtualization rely on high speed to deal with massive data throughput. Adopting these technologies more or less, most industries are facing inevitable data growth owing to the complex application scenarios. Correspondingly, the demand of the performance of infrastructure grows.

In light of such trend, more and more business and organizations adopt all-flash storage solution owing to the low latency, and the phenomenal performance in data rate makes all-flash devices suitable for various modern applications. The speed of SSD (solid state drive) depends on interface: While SAS (Serial Attached SCSI) and SATA (Serial Advanced Technology Attachment) remain common in the market, NVMe (Non-Volatile Memory Express) has achieved another breakthrough in speed and therefore gradually grows universal.

Recently receiving 5-star evaluation from the British media [ITPro](#), QSAN [XF3126D](#) is our most representative NVMe all-flash storage array product. QSAN XF3126D is equipped with Western Digital Ultrastar DC SN840 Gen3x4 PCIe NVMe enterprise-class dual-port SSD, which provides ultra-low latency in microseconds and therefore unprecedented speed, making the device the most competent and cost-effective flash storage in the market.

The current document mainly describes the performance of NVMe-based storage platform. Furthermore, an example in healthcare industry demonstrates how critical tasks could be accelerated with the deployment of all-flash storage.

The Evolution of NVMe Protocol


As SSDs gradually replaces traditional hard drives, SAS and SATA protocols are no longer sufficient to provide enterprises with the speed required to fully utilize their SSD storage. In light of this, NVMe, the latest drive interface, was developed specifically for non-volatile memory and is becoming more essential for data centers. The performance of NVMe SSDs outshines other drives, while the price per TB of storage is only slightly higher. Therefore, analysts such as IDC believe that NVMe will replace traditional storage protocols, especially for major workloads that are sensitive to latency. For more details about NVMe, please refer to the blog [What is NVMe and why is it important?](#)

Overview of NVMe SSD Speed

NVMe achieves a great leap on the speed by allowing SSDs to directly access the PCIe bus, reducing the latency to 30μ s (0.03ms). This helps the server to access directly connected NVMe SSDs. NVMe is able to do this because its command set requires less than half the number of CPU (Central Processing Unit) instructions to process I/O requests compared with the command sets of SCSI (Small Computer System Interface) and ATA (Advanced Technology Attachment). NVMe supports 64K commands in a message queue and up to 64K queues. Compare to the traditional protocol, SAS devices only support up to 256 commands per queue, while SATA supports up to 32 commands.

Pure NVMe Flash Storage

QSAN XF3126D, a 3U 26 bays all-NVMe flash storage, achieves the performance requirements of the enterprise high performance computing infrastructures with high IOPs at μ s-level latency. It's designed for enterprise users, providing excellent storage performance, enterprise-grade reliability, and a flexible and easy-to-use management system.



NVMe Dual-port Leadership

Western Digital Ultrastar DC SN840 is the 3rd generation of performance NVMe SSD for data center with PCIe Gen 3.1 (dual-port), NVMe 1.3, built on Western Digital's 96-layer 3D TLC NAND, with capacities up to 15.36TB in a U.2 2.5" form factor. The specification reaches up to 3,470/3,330 MB/s sequential read/write and up to 503K IOPS mixed random 70/30 read/write performance. In addition, it also provides 1 or 3 DW/D endurance classes and security options with encryption function.



2.5-inch U.2, 15mm, NVMe SSD
1.6TB, 3.2TB, 6.4TB,
1.92TB, 3.84TB, 7.68TB, 15.36TB¹

In comparison to Western Digital Ultrastar DC SN200, the previous generation of dual-port NVMe SSD, Ultrastar DC SN840 reaches a higher sequential write performance of up to 45%, and mixed IOPS (max, 70/30, 4KiB) up to 67%. In addition, SN840 saves 39% of power consumption while operating and 50% during idleness.

ITEM	WD DC SN840 1.92TB / 15.36TB	WD DC SN200 1.92TB / 7.68TB
Write Throughput (max MB/s, Seq)	2,280 / 3,190 ¹	2,100 ²
Mixed IOPS (max, 70/30 R/W, 4KiB)	231K / 401K	240K
Power (Operating)	18W	25W
Power (Idle)	6W	9W

¹ Seq 64KiB

² Seq 128KiB

QSAN NVMe AFA + WD NVMe SSD Brings Extraordinary Performance

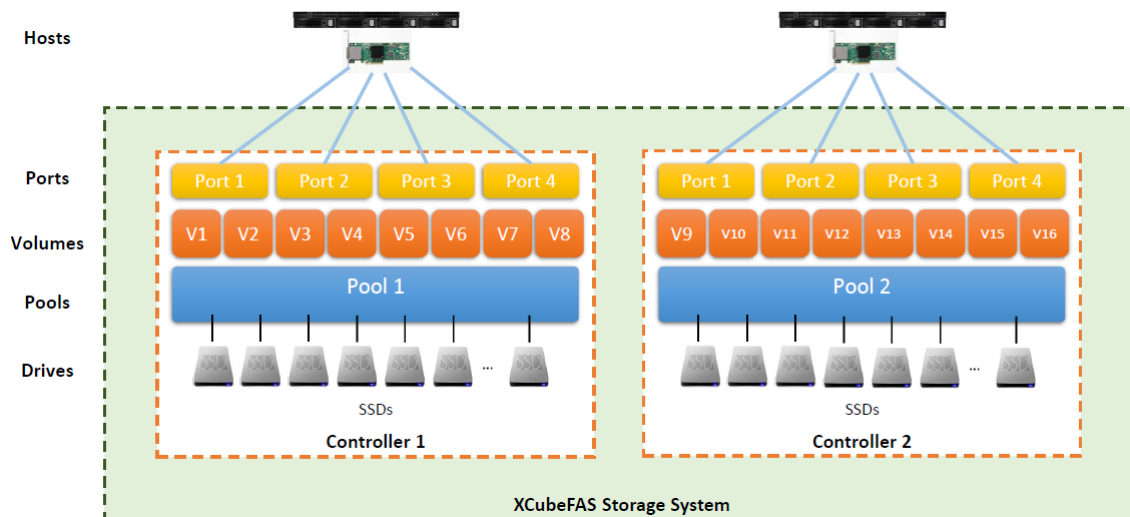
QSAN NVMe All-Flash Array plus Western Digital dual-port NVMe SSD provides a best practice performance results. Let's see how we do the test.

Test Equipment and Configurations

- Server
 - Model: 1 x Dell T630; 1 x HP Z840
 - 16Gb FC HBA: Marvell QLogic QLE2694
 - OS: Windows Server 2012 R2

- Storage
 - Model: XCubeFAS XF3126D
Memory: 16GB (2 x 8GB) per controller
Firmware 2.0.1
 - Dual-port NVMe SSD: 13 x Western Digital Ultrastar DC SN840, WUS4BA119DSP33Xz, 1.92TB, PCIe3.1 2x21
 - Pools:
 - 1 x RAID 5 / RAID 10 Pool with 7 / 6 x NVMe SSDs in Controller 1
 - 1 x RAID 5 / RAID 10 Pool with 6 x NVMe SSDs in Controller 2
 - Volumes:
 - 8 x 100GB in Pool 1 (Controller 1)
 - 8 x 101GB in Pool 2 (Controller 2)
 - Volume Stripe Size: 64KB
 - Block Size: 512 Byte
 - Cache Mode: Write-back in RAID 5, Write-through in RAID 10
- I/O Pattern
 - Tool: Vdbench V5.04.07
Workers: 2 x 4 (1 Worker to 1 Volume)
Outstanding (Queue Depth): 128
Xfersize: 4K 8K 32K 64K
I/O rates: 10 ~ 120
Reporting Interval: 1 sec
Warmup Period: 5 sec
Elapsed Time: 30 sec per I/O rate

Diagram




- Test Scenario

First we create two RAID 5 / RAID 10 pool. One for Controller 1 and the other is for Controller 2. Use the tool Vdbench and run I/O rates from 10% to 120%.

Performance Results

The charts below reveal the performance report. Under RAID 5 and RAID 6 configuration, the system respectively brings out up to 608K and 616K IOPS on random read; under the extremely low latency of less than 0.5ms, it still reflects the outstanding performance of nearly 230K IOPS.




On random write, both RAID 5 and RAID 6 achieve latency under 0.3ms with outstanding IOPS. It is worth highlighting that latency remains under 1ms while IOPS reaches more than 80K under both configurations; that is, the storage offers remarkable performance while parity bits are involved.

RANDOM WRITE 4K

RAID 5


152K IOPS



RANDOM WRITE 4K

RAID 6

137K IOPS




The performance of RAID 10 configuration is listed below for further comparison. The results reveal that the overall performance isn't seriously affected by parity bits. In light of this, the integration of XF3126D and SN840 is capable of providing phenomenal performance without sacrificing too much space for mirroring drives.

RANDOM READ 4K

RAID 10


639K IOPS



RANDOM WRITE 4K

RAID 10

615K IOPS




Application

The speed of NVMe-based storage platforms meets the requirement of technologies such as edge computing, artificial intelligence, and virtualization. Storage used to bottleneck certain computationally intensive tasks, and the NVMe offers solutions to the issue. For further understanding potential applications, the following paragraphs elaborate on an example of healthcare industry, showing how artificial intelligence and edge computing may together elevate the quality and efficiency of medical services.

The capability of machine learning is often limited by the efficiency of data input, during which data transmission and throughput become critical in order to generate more data within limited turnover time. Deploying all-flash storage devices in hospitals enables staff to employ artificial intelligence on, for example, automating resource scheduling and drug delivery. Moreover, potential applications vary from medical image interpretation to data analysis.

The power of artificial intelligence could be further promoted in collaboration with 5G network and edge computing, the key idea of which lies in transferring computing resources to the surrounding nodes of a network from its center. With accelerated speed on both storage devices and network, medical teams may provide better and more instant services even outside the hospital. For instance, edge computing may allow doctors and artificial intelligence perform treatment while confronting fatal circumstances on the ambulance. The topology below demonstrates a flash-integrated healthcare environment, which features more transparent information exchange and therefore improves the digitalization and quality of medical service.



Summary

The efficiency of computing tasks depends on data rate, and the speed of all-flash storage depends on drive interface. Designated for optimized performance, NVMe stands out among protocols in computationally intensive environments varying from machine learning to edge computing ecosystem.

The combination of QSAN all-flash array and Western Digital SSDs delivers a new generation of high-speed storage solution for organizations of all sizes, pushing industries forward on edge-cutting technologies development and integration. The low latency of all-flash devices contributes to the reliable environment to deal with massive data. Deploying flash storage does more than enabling organizations to conduct speed-demanding tasks. It paves the way to embracing the future trends.

References

- QSAN Product Page: XCubeFAS 3126D
<https://www.qsan.com/en/products/xf3126d>
- IT Pro: Qsan XCubeFAS XF3126D review: The price is right
<https://www.itpro.co.uk/infrastructure/server-storage/358833/qsan-xcubefas-xf3126d-review-the-price-is-right>
- QSAN Blog: Why Does Low Latency Affect Your Business
<https://blog.qsan.com/lower-latency-is-the-key-to-improve-competitiveness-for-enterprise/>
- QSAN News: QSAN Certifies Western Digital Ultrastar DC SN840
<https://www.qsan.com/en/newspage/qsan-certifies-western-digital-ultrastar-dc-sn840>
- Western Digital Document: Ultrastar SN840 datasheet
https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/product/data-center-drives/ultrastar-nvme-series/data-sheet-ultrastar-dc-sn840.pdf
- Western Digital Blog: What is NVMe and why is it important?
<https://blog.westerndigital.com/nvme-important-data-driven-businesses/>

- AnandTech: Best SSDs: April 2021
<https://www.anandtech.com/show/9799/best-ssds>
- Mobile Edge Computing Based QoS Optimization in Medical Healthcare Applications
https://www.researchgate.net/profile/Ali-Sodhro/publication/327876287_Mobile_edge_computing_based_QoS_optimization_in_medical_healthcare_applications/links/5bb2687d92851ca9ed33a9af/Mobile-edge-computing-based-QoS-optimization-in-medical-healthcare-applications.pdf