High Performance Enterprise SAN

XCubeSAN XS5226-D
Review Guideline

Part Number: QSWP1603A
Published: May 2016
Edition: 1.0

Edition 1.0 (May 2016)

This edition applies to QSAN XCubeSAN Series. Note that this document was produced based on beta code and some screens may change when it becomes generally available.

QSAN Technology, Inc.

4F., No.103, Ruihu St.,
Neihu Dist., Taipei City 114,
Taiwan (R.O.C.)

Tel: +886-2-7720-2118
Fax: +886-2-7720-0295

Email: sales@qsan.com
Website: www.qsan.com

Contents

Overview - Built for Work, Designed to Impress 1

Innovated Hardware Design & Stunning Performance 1

World's First 2U 26bay Enclosure 1
Advanced Cooling 2
Onboard 10GBase-T Ports 2
Dual Host Card Slots Per Controller 3
Green and Advanced C2F Memory Protection with Super Capacitor 3
Wake on SAS for Easy Management 4
Massive Scalability Up to 3PB 4

Functions to Try During Evaluation 5

High Availability 5
Amazing Performance by QSOE 2.0 Engine 5
SSD Optimized Storage 10

SANOS 4.0 : Enterprise Storage Functions 12

Data Protection & Backup 12
QThin (Thin Provisioning) 13
Fast Rebuild 13
QCache (SSD Caching) 13
iSCSI Force Field for Data Security 13

Service, Support & More Information 13

Product Support 13
Purchasing Information 13

Overview - Built for Work, Designed to Impress

The most advanced modern technology meets the most outstanding QSAN storage yet. The result of this combination? A whole new era of power, performance and capacity. XCubeSAN XS5200 series gives all the business customers the performance and reliable storage you need. Now with the newest generation Intel Xeon processor, QSAN thoughtfully created the XS5200 series for work intensive business critical applications.

XS5200 series proudly supports a complete range of enclosure form factors from 2U 12bay, 2U 26bay to 3U 16bay and 4U 24bay. Especially 2U 26bay model – XS5226-D is the world's highest density SAN in 2U form factor. 26 bays can be configured as 24+2 combination, 24 stands for 24 bays for data capacity and 2 stands for 2 bays for SSD caching use. It will efficiently improve the performance of system.

XCubeSAN XS5226

Innovated Hardware Design & Stunning Performance

In order to handle the volume and velocity of the ever increasing data, XS5200 series is built on the next generation storage platform with the upmost quality. Hardware innovations include the following.

World's First 2U 26bay Enclosure

More durable than its predecessor, and part of the most durable QSAN XS5200 SAN family, the new XS5226-D undergoes extensive testing to ensure it can survive business customers' everyday workday, and features premium materials like enhanced HDD tray and HDD door; which make QSAN can be the global first one to create the 2U 26 bays storage system.

QSAN takes a new approach to use Talc in plastic compounds, which will stiffen thermoplastics, increase thermal conductivity & deflection temperature, enhance the strength; and our advanced mechanical technology comes with Glass Fiber Reinforced Plastic to create the new durability in solid structure design.

The honeycomb pattern on the grille of HDD door not only reduces wind shear noise but also channels the airflow into more focused beam.

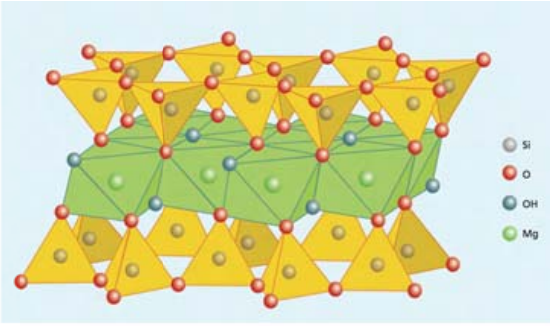

Advanced Cooling

High density system design always comes with thermal issues. XS5226-D solves thermal issues not only by heat dissipation mechanism, but also airflow of cooling fan and chassis fine tune. Heat dissipation mechanism links main heat sources and conduct heat outside the enclosure to keep components within operation temperature.

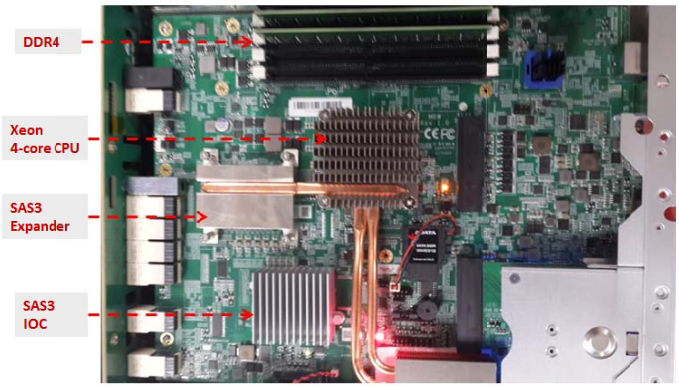

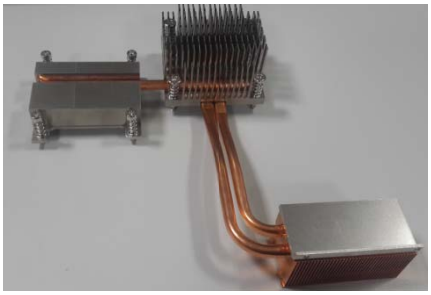

Onboard 10GBase-T Ports

XS5226-D base system is equipped with onboard 4x 10G BASE-T ports. Without any extra host cards for extended host server connection, XS5226-D can already provide 40Gb iSCSI bandwidth to meet most IT and datacenter requirements for 10GbE solutions.

Dual Host Card Slots Per Controller

With fault tolerance and modular redundancy design in mind, XS5200 has 4 slots to expand host connectivity. Three types of host card are available for selection: 4-port 16Gb FC host card, 4-port 10GbE host card and 4-port 1GbE host card. XS5226-D supports both iSCSI and Fibre Channel protocols at the same time. You may start the initial base unit controller without any host card installed. As your business grows, you can purchase one of the three host cards to easily acquire extra bandwidth to the servers. Because there are so many host connection ports available, you can save the money for purchasing Fibre Channel switch or Ethernet switch. This is a great way to start your investment small and grow bigger as your business grows larger.

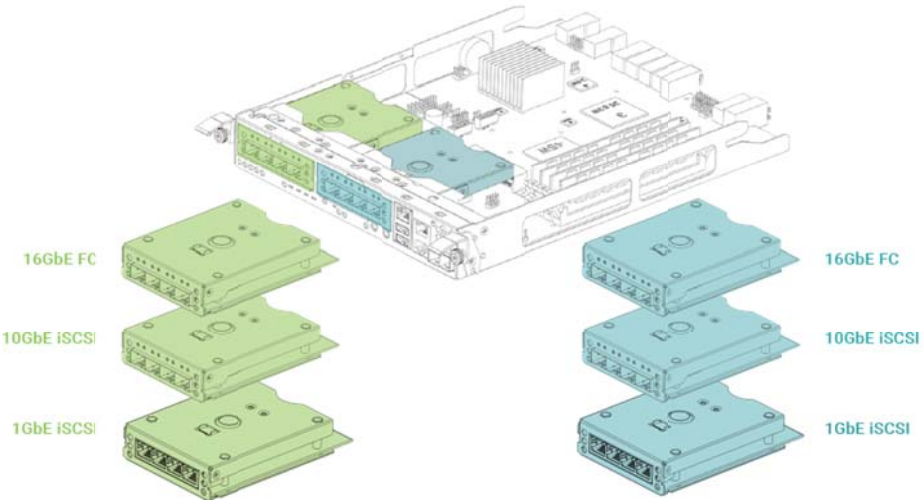

Green and Advanced C2F Memory Protection with Super Capacitor

Cache memory protection is a must-have in enterprise storage systems. In XS5226-D, C2F (Cache to Flash) memory protection solution is composed of C2F power module and C2F flash module. C2F power module has two types: SCM (Super Capacitor Module) and BBM (Battery Backup Module).Both flash module and power module are hot pluggable from the back of the enclosure just like the SAN controller. This design makes maintenance super easy and no downtime because you don't need to plug out the SAN controller to service the flash and battery modules. To raise the bar even higher, the flash module adopts M.2 PCIE interface to provide the fastest and safest storage to store cache data when unexpected power loss happens during operation. Super capacitor has several benefits over BBM. They are green (no heavy metals and no disposal issues) and environment-friendly, low maintenance and long lifetime, and air transport-friendly.

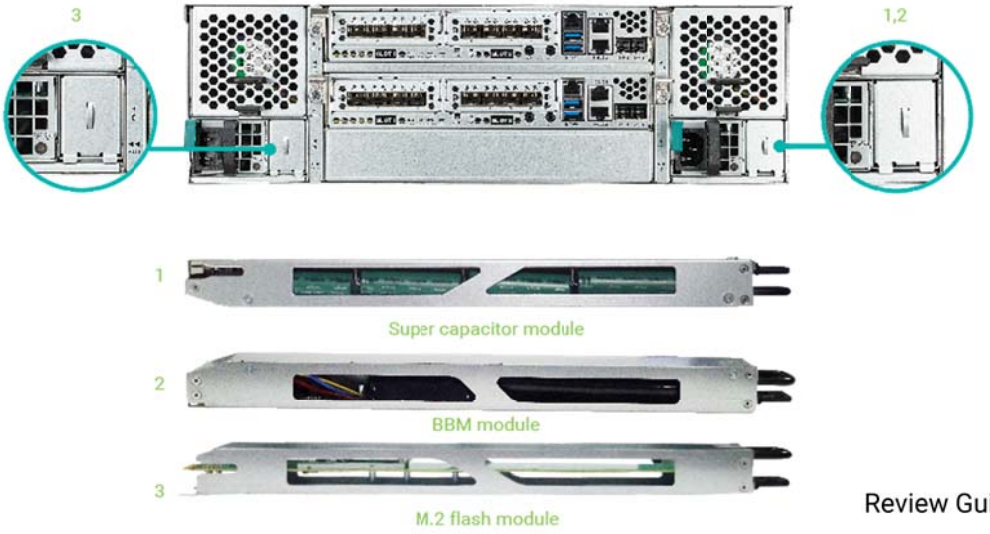

Wake on SAS for Easy Management

XS5200 series provides wake on SAS function that can power on/off all cascaded XD5300 expansion enclosures by using QSAN proprietary SAS cable. The expansion enclosure will synchronize system power the same as XS5200. It reduces remote maintenance effort and power consumption.

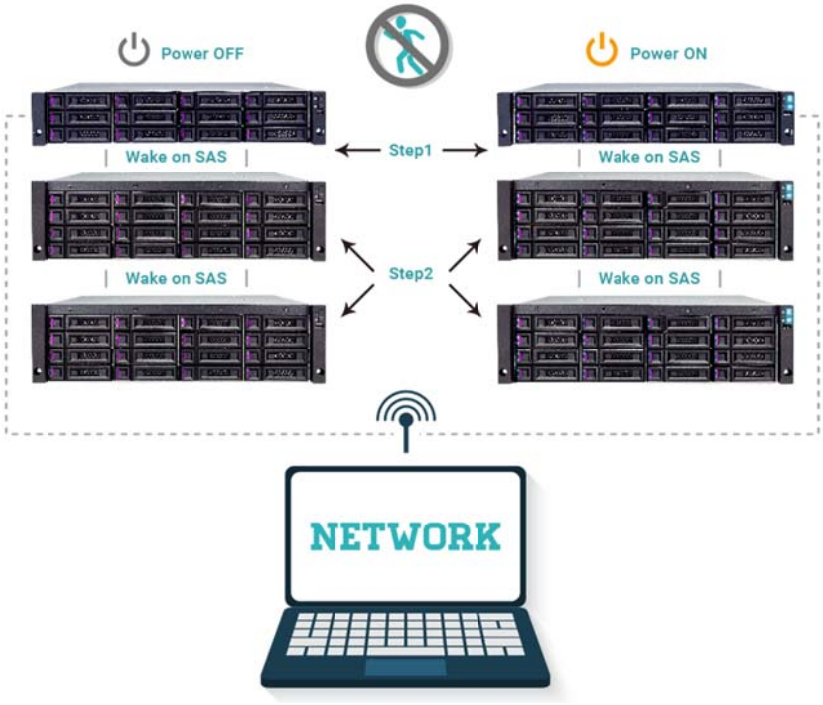

Massive Scalability Up to 3PB

Each SAN controller is equipped with 2x mini SAS HD (SFF-8644) expansion ports with 96Gb/s bandwidth to guarantee that aggregated disk I/O can pass on to SAN controller without traffic jam. XS5226-D can expand capacity by cascading disk enclosures – XD5300 series from XCubeDAS product line. Connecting up to 10 expansion enclosures makes XS5226-D capable of delivering 3PB capacity using 8TB nearline SAS drives.

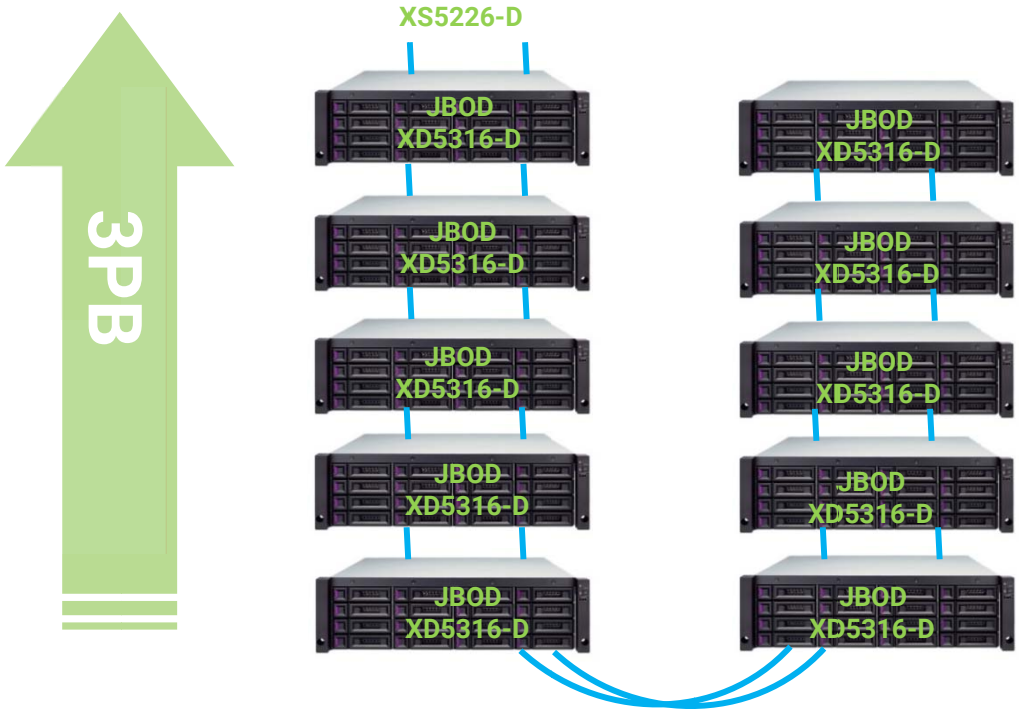

Functions to Try During Evaluation

XS5226-D model is powered by SANOS 4.0, which is the SAN operating system of all QSAN SAN storage systems. SANOS 4.0 is the heart and soul of XS5226-D model and made of 64 bit embedded Linux with proprietary enterprise storage functions including the following:

High Availability

XS5226-D is built upon the next generation storage platform made of redundant and fault tolerant modules and can provide 99.999% high availability in many ways. All these redundant modules are hot pluggable to provide non-stop service.

- Dual active SAN controllers : Support ALUA
- Multiple I/O path and load balance : Support MPIO and MCS
- Redundant cooling fan module, power supply module, and HDD trays
- Non-disruptive firmware upgrade
- RAID protection with parity : Support RAID 5, 6, 50, 60
- Automatic management service failover

XS5226-D has two active-active SAN controllers. The memory cache of each controller is mirrored together through PCI-E channels on the midplane to synchronize with each other. Both controllers can deliver storage services (host I/O, snapshot, replication, cloning) at the same time. In case of controller failure, all the services will automatically failover to the peer controller and the failover process can complete in less than 10 seconds.

With two host cards per controller, XS5226-D can yield up to 20x 10GbE iSCSI ports or 16x 16Gb Fibre Channel ports. Multiple I/O path function is built-in by default to aggregate these I/O ports into higher bandwidth and provide load balancing function. This is crucial for clustered server application and virtualized datacenters. Thanks to dual active controller design, when you perform SAN controller firmware upgrade, the two controllers can take turns to reboot to achieve non-disruptive firmware upgrade.

The philosophy of SANOS 4.0 is all about making IT manager's life more efficient and more productive. All modules are hot pluggable for easy replacement in the field. When the master controller fails, the management port will automatically failover to the other controller as well with the same IP address. So IT managers can remain in full control without physical attendance next to the storage enclosure.

Amazing Performance by QSOE 2.0 Engine

QSAN Storage Offload Engine (QSOE 2.0) can offload the overheads of processing iSCSI and Fibre Channel protocols to the dedicated cores of Xeon CPU. The result is 50% to 60% overall performance improvement in iSCSI and Fibre Channel.

With 16Gb FC card in slot 1 and 10GbE iSCSI card in slot2, XS5226-D can deliver **12,480MB/s** sequential read and **8,000MB/s** sequential write in throughput and **1,950,000 IOPS** in sequential.
Below is the recommended procedure for testing performance.

Storage system requirements

- One XS5226-D dual controller system with 16Gb FC card in slot 1 and 10GbE iSCSI card is slot2
- 26x 12Gb SAS SSD drives (read throughput over 800MB/s)

Host server requirements

- At least four powerful host servers
- Intel Xeon E5-2620 @ 2GHz or higher
- At least 24GB RAM
- 16Gb FC HBA with at least 2x ports (SFP+)
- 10GbE iSCSI HBA with at least 2x ports (SFP+)
- Windows Server 2012 R2
- Benchmark tool : IOmeter

Architecture scheme

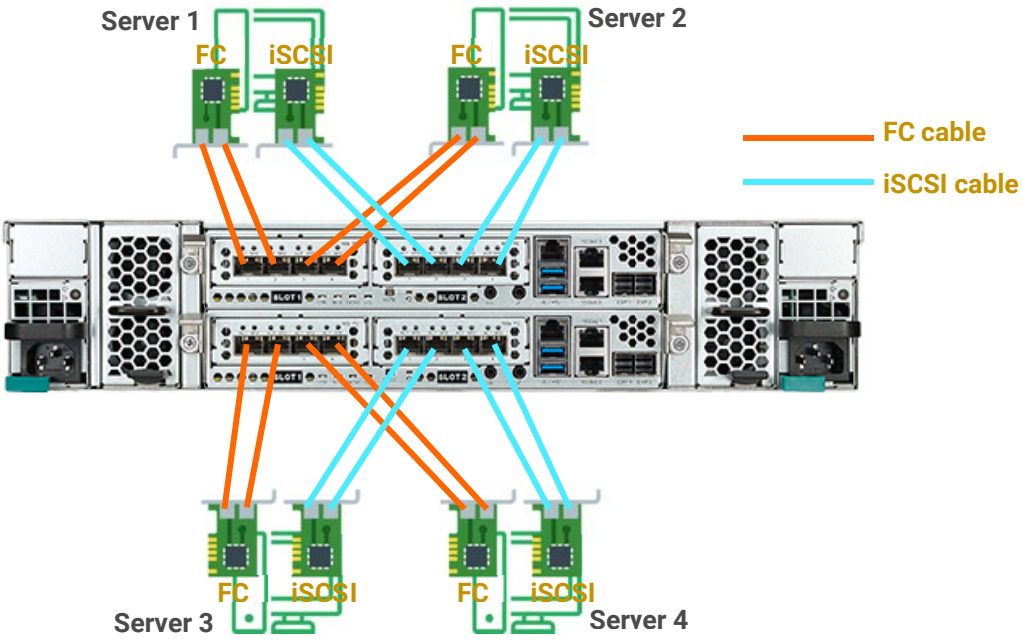

Storage pool configuration for IOPS performance test

Each pool accessed by FC connection is composed of two SAS SSDs, while accessed by iSCSI connection is composed of one SAS SSD. Each pool provides one volume and mapped to one FC port or 10GbE port. Totally, 24x SAS SSD drives are involved.

- 8 pools: RAID 0, 2 SSDs per pool, one 16Gb FC port to access one pool
- 8 pools: RAID 0, 1 SSD per pool, one 10GbE port to access one pool

Storage pool configuration for throughput performance test

Each pool accessed by FC connection is composed of three SAS SSDs. Each pool provides one volume and mapped to one FC port. Totally, 24x SAS SSD drives are involved.

- 8 pools: RAID 0, 3 SSDs per pool, one 16Gb FC port to access one pool

Settings to get best IOPS and throughput

Please make sure that FC HBA registry setting's queue depth is set to 255. Take QLogic FC HBA registry setting for example in the following screenshot.

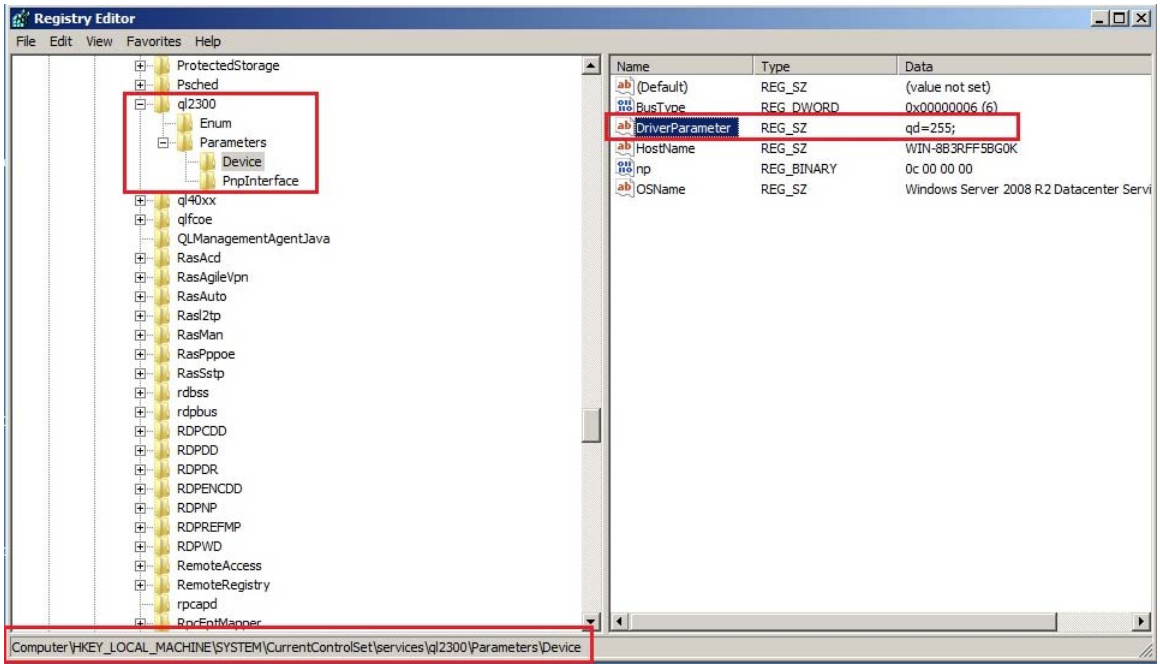

In order to get the best IOPS, please change settings on FC HBA registry and Iometer. Iometer setting: outstanding I/O set to 256, transfer request size set to 512B.

In order to get the best throughput, please set Iometer settings: transfer request size to 1MByte and outstanding I/O to 16.

Pool and volume creation

Please follow below procedures to setup your storage pools and volumes for performance evaluation.

1. **Step1**. Click **Pools** under **STORAGE MANAGEMENT**, then click **Create Pool** icon to open Create Pool window.

   Assign Pool Type to Thick Provisioning. Enter Pool Name for the pool and maximum length is 16 characters. Select Preferred Controller from the drop-down list, I/O resources will be managed by the preferred controller; this option is only available on dual controllers system. Press Next button when all settings are done.

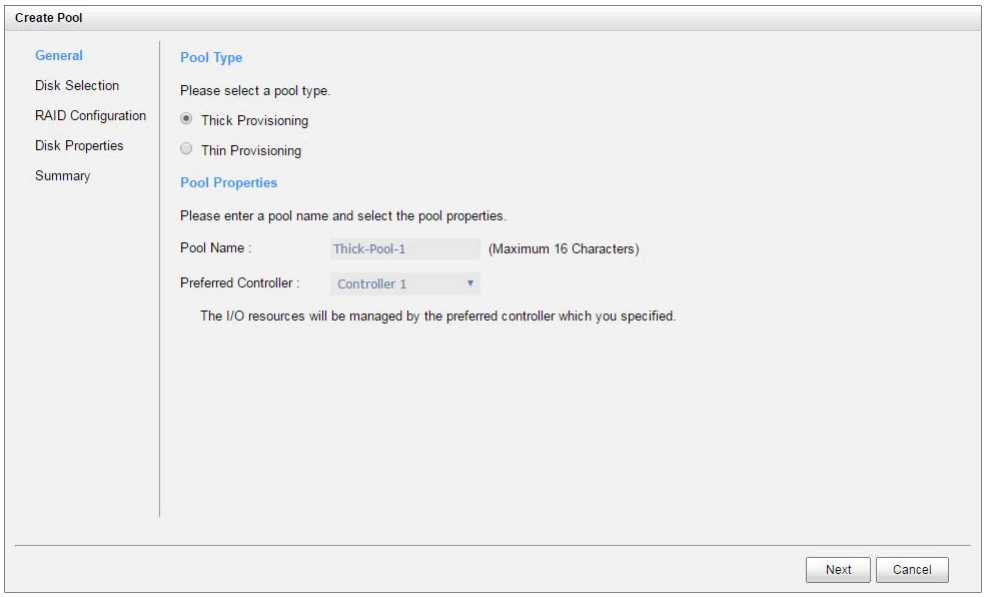

2. **Step2**. Select 2x SAS SSD drives for thin provisioning pool. Check 2x SAS SSD drives from the drive table for new pool use, then press **Next** button.
3. **Step3**. Select RAID level to RAID 0 for pool, then press **Next** button.
4. **Step4**. Set Disk Properties if necessary. When disk properties set is done, press **Next** button to get summary of new pool.

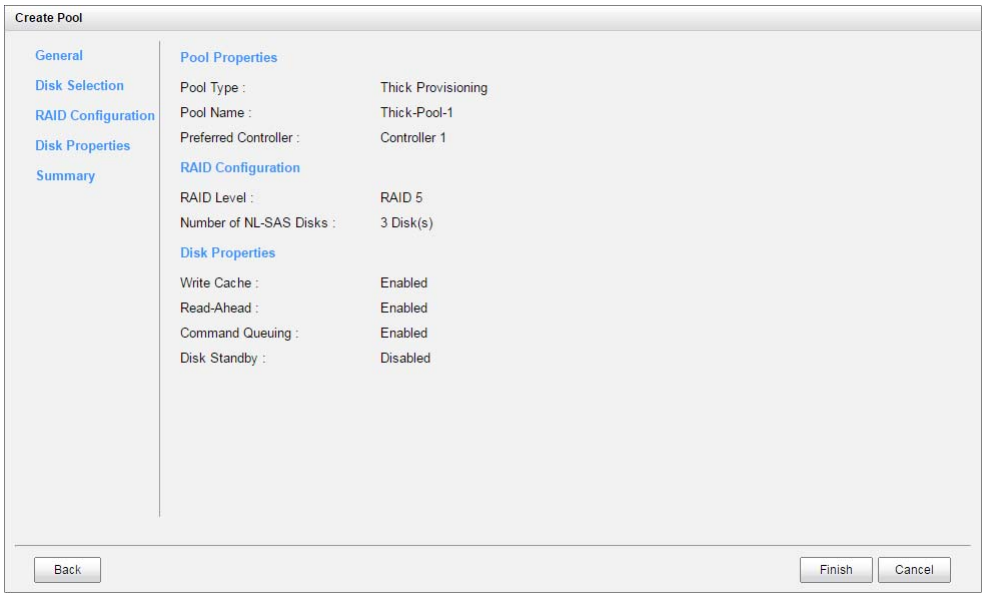

5. **Step5**. Create a volume that allocated from above new pool. Click **Volumes** under **STORAGE MANAGEMENT**, then click **Create Volume** icon to open Create Volume window.

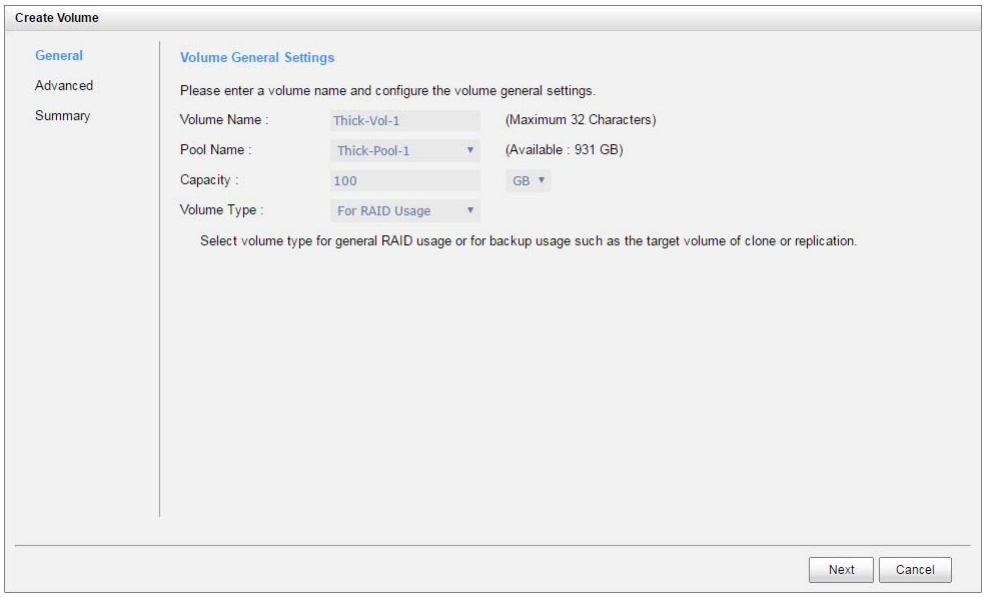

6. **Step6**. Volume advance settings by default settings. Press **Next** button.
7. **Step7**. Summary of new volume are shown as below. Press **Finish** button.

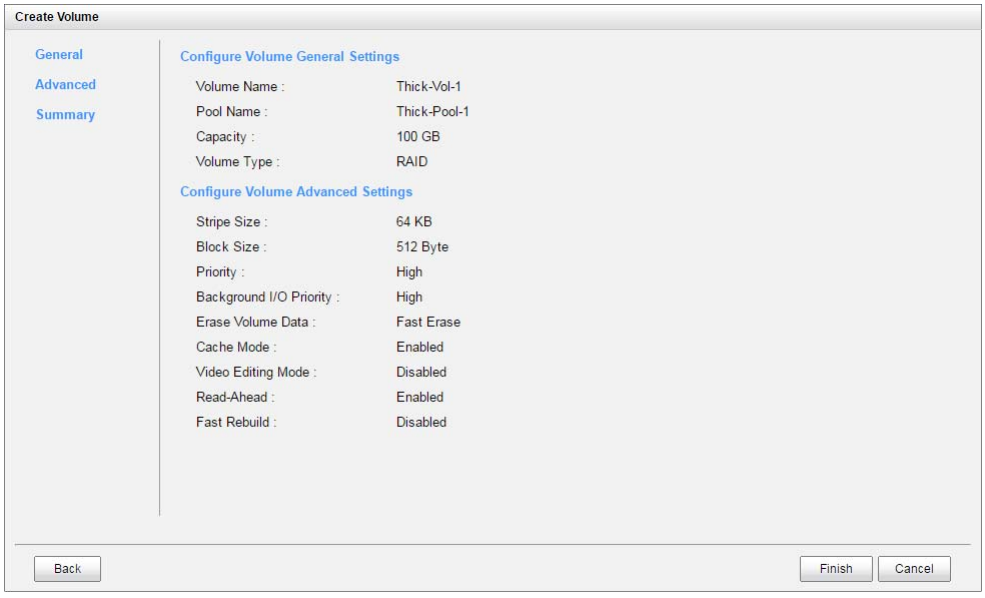

IOmeter Configuration

- Version 2006.07.27 Dynamo
- Queue Depth (QD, outstanding I/Os):

  set to 256 for IOPS performance test

  set to 16 for throughput performance test
- Test cases

100% Read; Sequential for IOPS: 512B, 4KB, 1MB

100% Read; Sequential for Throughput: 512B, 4KB, 1MB

100% Write; Sequential for IOPS: 512B, 4KB, 1MB

100% Write; Sequential for Throughput: 512B, 4KB, 1MB

SSD Optimized Storage

The latest inclusion in SANOS 4.0 is QTiering (Auto Tiering) function. Similar to QCache (SSD Caching) function, QTiering leverages the benefits of SSD drives to enhance overall system performance by allocating volumes to different storage tiers according to data access frequency (or data hotness). The benefit is lower TCO and saving money in provisioning overall storage capacity while maintaining appropriate levels of performance.

There are two major differences between QCache and QTiering. The first one is that the size of SSD cache is not part of the storage capacity. When you shut down the storage system, everything in SSD cache is gone. However, the size of SSD tier is part of the size of the storage pool. When you shut down the storage system, everything in SSD tier remains intact. The second one is that QCache can improve random I/O performance, especially random read. QCache doesn't help too much in throughput or sequential I/O. However, QTiering can enhance both sequential and random I/O greatly at the same time.

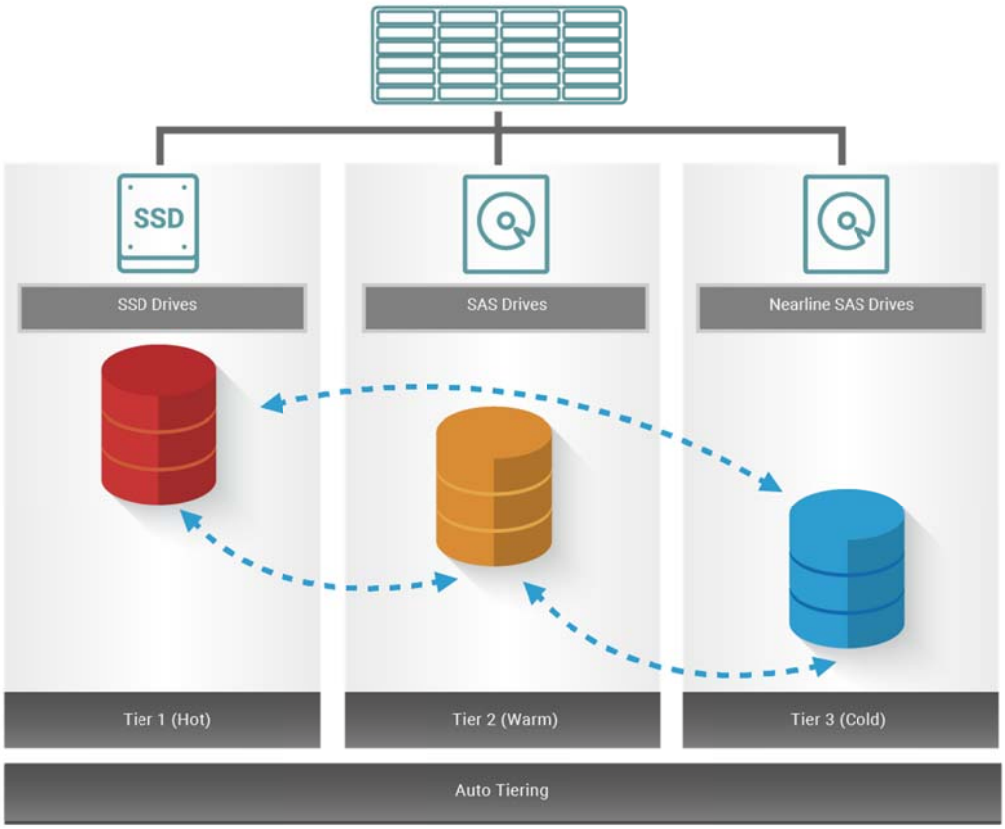

Highlights of QTiering

QTiering is easy to use and has intuitive web management interface. Some highlights of QTiering :

- Support 4 storage tiers :

  Tier 1: SSD drive

  Tier 2: 15K/10K SAS drive

  Tier 3: 7.2K nearline SAS

  Tier 4: SATA drive (only in single controller mode or with 2U 26 enclosures)

- Dynamic volume tiering policies, Each LUN has its own policy. It can be changed anytime according to application needs.

  Policy 1 : Highest available tier

  Policy 2 : Auto tiering

  Policy 3 : Start highest then auto tiering

  Policy 4 : Lowest tier

  Policy 5 : No data movement

- Intelligent algorithm manages data relocation according to tiering policies and the ranking results. The LUN volume is divided into units of 1GB size, which is called sub LUN. Access records to each sub LUN are collected and analyzed every hour. Intelligent algorithm uses these statistics to calculate the data hotness ratio using half-life coefficient and advanced ranking mathematics to generate the ranking results.

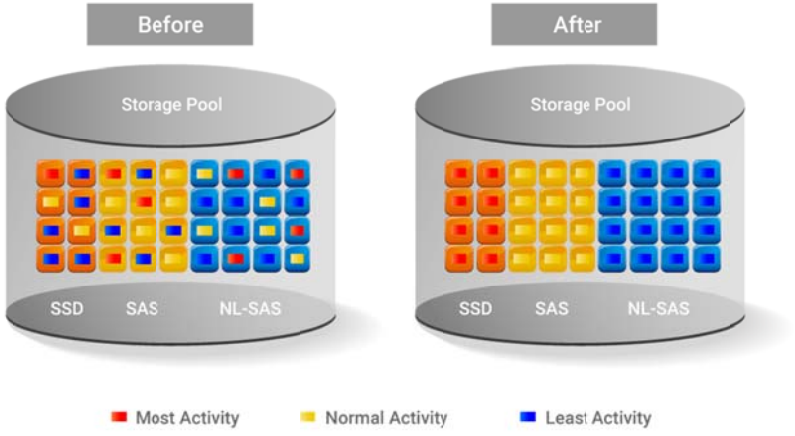

- When relocation is in process, there is no downtime. Users can enjoy the performance benefits brought by QTiering. As the volume data is gradually moved to the faster tier, the performance will improve accordingly. The following table shows the performance benefits.

Transfer speed comparison for file copy

Data	Data tier	On-going data relocation	Transfer speed (MB/s)
16GB zip file	Tier 1	no	386MB/s (*1)
16GB zip file	Tier 1 and Tier3	yes	291MB/s (*2)
16GB zip file	Tier 3	no	275MB/s (*3)

**1. Data in Tier 1 (SSD), took 45 seconds (386 MB/s)*
**2. Data relocating from Tier 1 to Tier 3, takes 58 seconds (291 MB/s)*
**3. Data in Tier 3 (NL-SAS), took 64 seconds (275 MB/s)*

SANOS 4.0 : Enterprise Storage Functions

Data Protection & Backup

XS5226-D provides QSnap (snapshot), QClone (local cloning), and QReplica (remote replication) functions to safeguard your business data. All three functions support scheduled tasks. The best part is they are all free and supported by default.

QSnap adopts copy-on-write technology to keep a point-in-time record of the difference of volume data. It's the most efficient way to make data backup. Especially when ransomware attack is outrageous these days, QSnap can help you quickly recover your maliciously-encrypted data and set you free from being blackmailed.

- Support system-wise 4096 snapshots
- 64 snapshots per LUN/volume
- 64 LUNs/volumes can be enabled with snapshot

QReplica prevents primary site failure by replicating data to the remote sites. It is easy to setup on both sites over Ethernet. Its function is based on QSnap technology to replicate incremental data changes. The replication could be scheduled to prevent replication during office hours.

- 1 replication task per volume
- 32 replication tasks per system
- 8 task shaping settings

QThin (Thin Provisioning)

Please check https://www.qsan.com/en/software.php?no=71D34464

Fast Rebuild

Please check https://www.qsan.com/en/software.php?no=71D34464

QCache (SSD Caching)

Please check https://www.qsan.com/en/software.php?no=71D34464 or download the white paper at https://www.qsan.com/data/editor/files/QWP201402-SSD_Caching_v1_3.pdf

iSCSI Force Field for Data Security

Please check https://www.qsan.com/en/software.php?no=71D34464

Service, Support & More Information

Product Support

XS5200 series including XS5226-D is covered by a 3 year limited warranty. Technical support and related downloads are offered via QSAN official website at http://www.qsan.com

Purchasing Information

XCubeSAN XS5200 series products are available at popular resellers and distributors worldwide. For more information, please visit “Where to Buy” section in http://www.qsan.com

For assistance during your evaluation, please contact:

marketing@qsan.com

QSAN Technology, Inc.

+886 (0)2 7720-2118 ext 166