



Data Storage, Security & Performance

XCubeSAN

| XS5200 Series



Product Highlights

- High Performance SAN storage system with Dual-Active (Active/Active) controller
- High availability design with no single point of failure
- 5th generation Intel® Xeon® 4-core processor, up to 128GB RAM per controller
- Latest 12Gb SAS 3.0 technology
- Built-in 10GbE iSCSI
- Up to 12,000MB/s sequential read and 8,000MB/s sequential write throughput, up to 1.5 million sequential IOPS
- Scale up solution supports over 8.7PB of raw storage capacity
- QSAN SANOS (SAN Operating System) v4.0
- Advanced Storage Management
 - Thin Provisioning
 - SSD Cache (read and write cache)
 - Auto Tiering
 - Snapshot
- Flexible I/O host cards for iSCSI SAN or Fibre Channel SAN
- Local clone and remote replication for disaster recovery
- Virtualization support for VMware VAAI, Microsoft Hyper-V ODX, and Citrix
- Cache-to-Flash memory protection technology





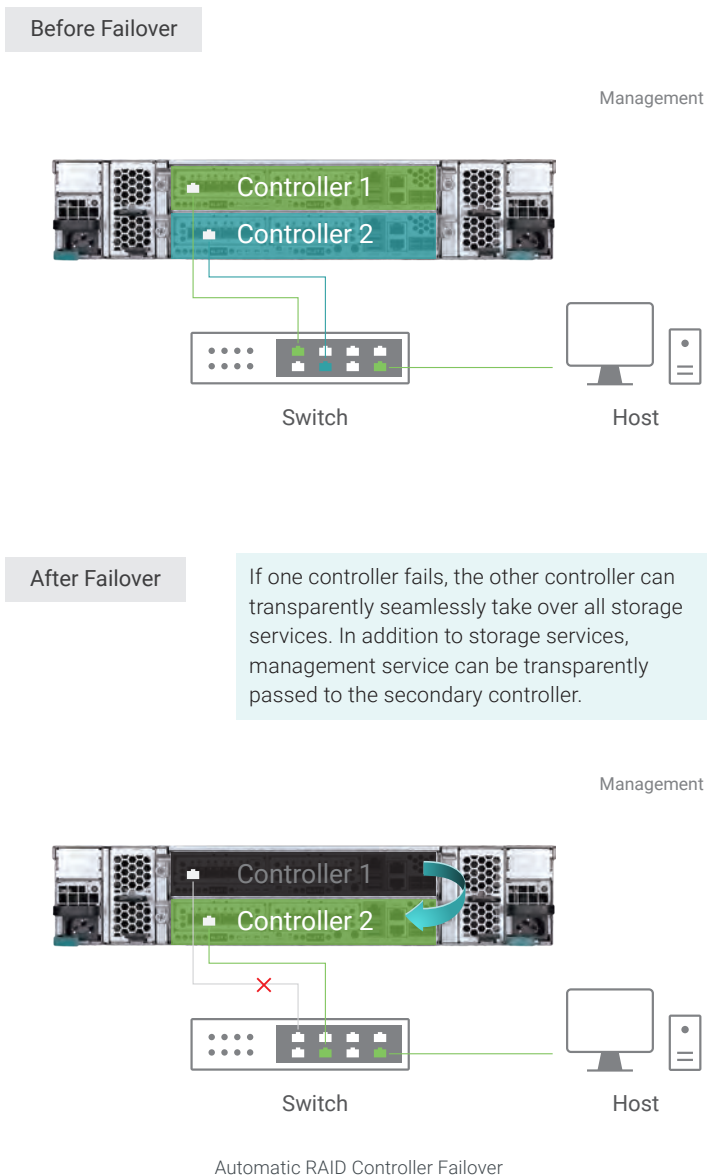
- [P1 / Table of Contents](#)
- [P2 / Overview](#)
- [P2 / Dual-Active \(Active/Active\) Controller SAN System](#)
- [P3 / High Availability with No Single Point of Failures](#)
- [P5 / Ultra High Performance](#)
- [P7 / Wide Ranging Product Portfolio](#)
- [P7 / QSAN SANOS 4.0 SAN Operating System](#)
- [P8 / Advanced Storage Management](#)
- [P10 / Thin Provisioning \(QThin\)](#)
- [P12 / SSD Cache \(QCache 2.0\)](#)
- [P13 / Auto Tiering \(QTiering\)](#)
- [P14 / Snapshot \(QSnap\)](#)
- [P15 / High Reliability](#)
- [P15 / Flexible Scale-up Solution](#)
- [P17 / Modular I/O Ports for iSCSI SAN and Fibre Channel SAN](#)
- [P18 / Cache-to-Flash Memory Protection Technology](#)
- [P19 / Comprehensive Data Security](#)
- [P20 / Data Backup & Disaster Recovery](#)
- [P24 / Virtualization](#)
- [P25 / Interoperability](#)
- [P25 / Wide Compatibility](#)
- [P25 / Optional USB LCM](#)
- [P26 / Green Technology](#)
- [P28 / XCubeSAN Applications](#)
- [P30 / Accessories](#)
- [P31 / Hardware Specifications](#)
- [P32 / Software](#)

XCubeSAN XS5200 Series Overview

The XCubeSAN XS5200 series is the newest generation of Dual-Active (Active-Active) SAN systems with high performance. It's simple, secure, scalable, and affordable for enterprises. It brings enterprise-level features such as thin provisioning, SSD read and write cache, auto tiering, and snapshot to the SMBs. The XS5200 products are designed to be an ideal solution to the applications of mission critical data centers, high performance computing, virtualization integration, or media and entertainment environments.

Dual-Active (Active/Active) Controller SAN System

The XS5200 series features a Dual-Active controller architecture. Both controllers concurrently provide storage services in real time. Active-Active architecture doubles the available host bandwidth and cache-hit ratio which ensures there is no wasted resource in the system.



High Availability with No Single Point-of-Failure

The XS5200 series is a proven high-availability SAN storage system. All of the critical components in the XS5200 series are hot pluggable and designed with full redundancy. Additionally, dual-active controller design and automatic failover/failback mechanism and cache mirroring through NTB (Non-Transparent Bridge) bus to achieve Active-Active HA functionality allows the XS5200 series to withstand multiple component failures and achieve 99.999% availability.

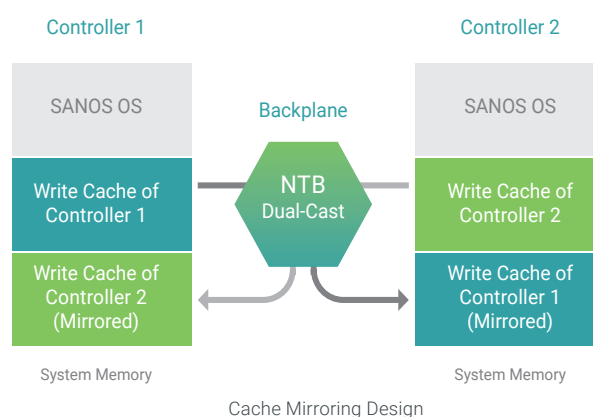
• Dual-Active Controllers

Active-Active design and automatic failover/failback mechanism provides the highest level of service availability and supports non-disruptive firmware upgrades. Dual-Active design also allows double the host bandwidth and cache-hit ratio, utilizing both controllers also means there are no idle resources within the system.

• Cache Mirroring through NTB Bus

The DDR4 ECC system memory in the XS5200 SAN controller is used by the SANOS and I/O cache. In order to achieve Active-Active HA functionality, the write cache on both SAN controllers needs to be identical and synchronized in real time. When one controller fails, the other controller can seamlessly take over all the tasks of the failed controller.

The XS5200 series achieves this by leveraging the NTB (Non-Transparent Bridge) hardware within the Intel® processor allowing for full failover protection.



• Redundant and Hot-pluggable Components

The XS5200 series features a fully modularized, cable-less architecture. In addition to Dual-Active controllers, all critical components inside the unit such as power supply modules, fan modules, and hot pluggable to provide fault tolerance capabilities.

In case of any component failure, the system will notify the administrator immediately; an alarm will go off and the web management interface will clearly indicate which component is at risk or has failed. The IT manager can then simply unplug the failed component and replace it without affecting uptime.



Fully Redundant Modular Design

RAID Protection and Redundant 12Gb SAS Expansion

The XS5200 series drive trays are dual-ported and can accept I/O requests from both SAN controllers, providing redundant I/O paths.

SANOS 4.0 supports 15 different RAID levels to provide redundancy in the storage pool for an additional layer of protection. The global hot spare function enhances RAID protection by automatically replacing the failed disks and starting the rebuild process without need for user intervention.

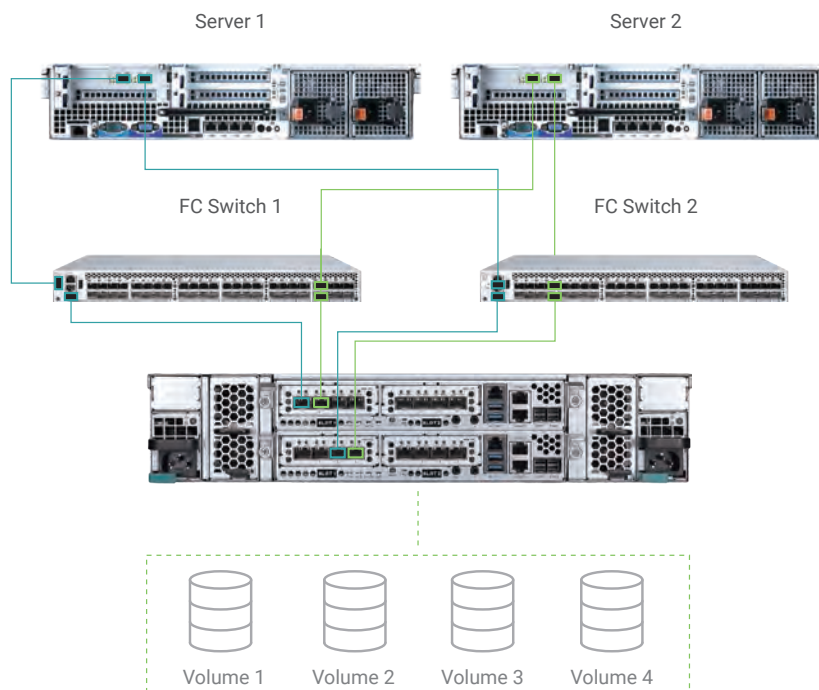
There are two dedicated mini-SAS HD (SFF-8644) ports on each controller to provide path redundancies to XD5300 series expansion enclosures. All these redundancies help the XS5200 series achieve 99.999% availability.



Two 12Gb SAS Expansion Ports per Controller

Multipath Storage Access

MPIO (Multi Path Input Output) is a fault-tolerance and performance enhancement technique which allows the use of more than one path to the volume from the same host server. MPIO is supported in both iSCSI and Fibre Channel protocols.

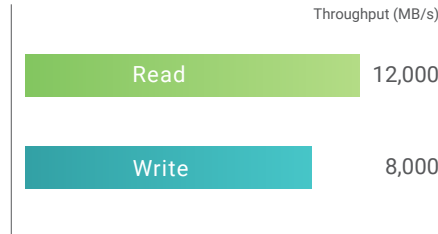


MPIO High Availability Topology for Clustered Servers

The XS5200 series supports the ALUA (Asymmetric Logic Unit Access) standard. ALUA uses SCSI 3 primary commands that are part of the standard SCSI SPC-3 specification to provide alternative I/O path capability to protect against port failure. With ALUA support, I/O of the same volume can be sent through either SAN controller. Not only can MPIO provide path redundancy for high availability, it also improves and scales up performance.

Ultra High Performance

With its all-new hardware architecture and leveraging Intel® processor, 12Gb SAS 3.0 backplane, built-in 10GbE LAN, and the finely-tuned SANOS 4.0's performance, the XS5200 series can deliver an astounding 12,000MB/s¹ sequential read and 8,000MB/s¹ sequential write in throughput and over 1.5 million² sequential IOPS.



• Latest 12Gb SAS 3.0 Controller

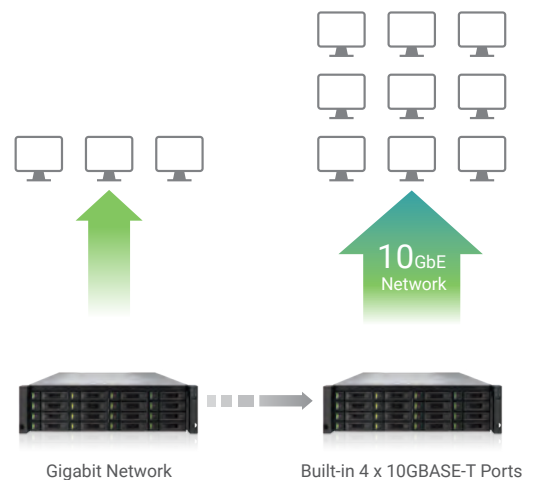
12Gb SAS 3.0 is the latest and fastest drive interface available. It doubles the data transfer rate of SAS 2.0 (6Gb), while remaining backward compatible with 6Gb SAS 2.0 drives. The benefit is that you have the flexibility to purchase less expensive 6Gb SAS drives for initial installation or leverage investment of your existing 6Gb SAS drives and have a peace of mind knowing you can migrate to 12Gb SAS 3.0 by purchasing only new drives.



12Gb SAS is backward compatible with 6Gb SAS

• Built-in High Speed 10GbE LAN Ports

The base unit SAN controller comes with two onboard 10GBASE-T iSCSI ports per controller. The Dual-Active controller XS5200 SAN system has four 10GBASE-T iSCSI ports reaching 40Gb/s bandwidth when you aggregate all ports together. This base unit SAN system comes ready to support a variety of applications including data sharing, backup, video editing, and native virtualization support for VMware, Citrix, and Hyper-V.




¹ Benchmarked with 16GB Memory (2 x 8GB) in two different banks, IOmeter utility, 512KB I/O size, non-cache hit, 128 queue depths, and 26 x 12Gb SAS SSD drives.

² The number is derived from sequential, non-cache hit, small I/O size (4KB) using 12Gb SAS SSD drives.

· 4-port 16Gb Fibre Channel Host Card

The XS5200 series supports the 4-port 16Gb Fibre Channel host card, it can provide as many as 8 x 16Gb FC ports for a staggering 128Gb/s super wide host server bandwidth. The benefit of 16Gb FC ports is greater connectivity to more host servers, this will allow you to consolidate your storage from multiple SAN's with only four ports, into one system with eight ports.

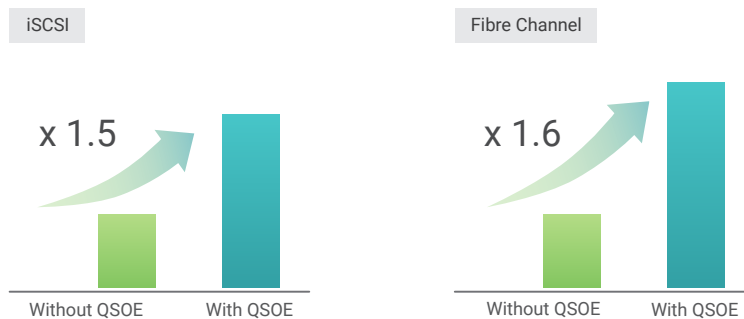
16Gb Fibre Channel is capable of twice the data transmission speed of older 8Gb Fibre Channel system and is ideal for high-end applications such as live databases management and virtualized datacenters. Fibre Channel also reduces the number of connection cables, HBA cards, and reduce power consumption by increasing the workload of a single server and expanding the number and size of concurrent applications instead of using multiple 8Gb Fibre Channel SAN systems to achieve the same performance.



Up to 8 x 16Gb Fibre Channel Ports

· QSOE (QSAN Storage Optimization Engine)

QSOE, one of SANOS 4.0 software modules, can optimize communication centric processes to reduce protocol overheads, increase session scalability and therefore increase total I/O throughput. As a result, iSCSI throughput can be boosted up to 1.5 times and Fibre Channel throughput up to 1.6 times.



Boosting Performance by QSOE Technology

Wide Ranging Product Portfolio

The XS5200 series features a wide range of form factors including a 24-bay, 4U 3.5" LFF chassis (XS5224), 3U 16-bay (XS5216), 2U 12-bay (XS5212), and a 26-bay, 2U 2.5" SFF chassis (XS5226). This allows more deployment flexibility to meet users' budgets and rack density limitations. Through the modular I/O expansion slot, the XS5200 series models can be easily configured as iSCSI SAN storage or Fibre Channel SAN storage or even a hybrid of the two.

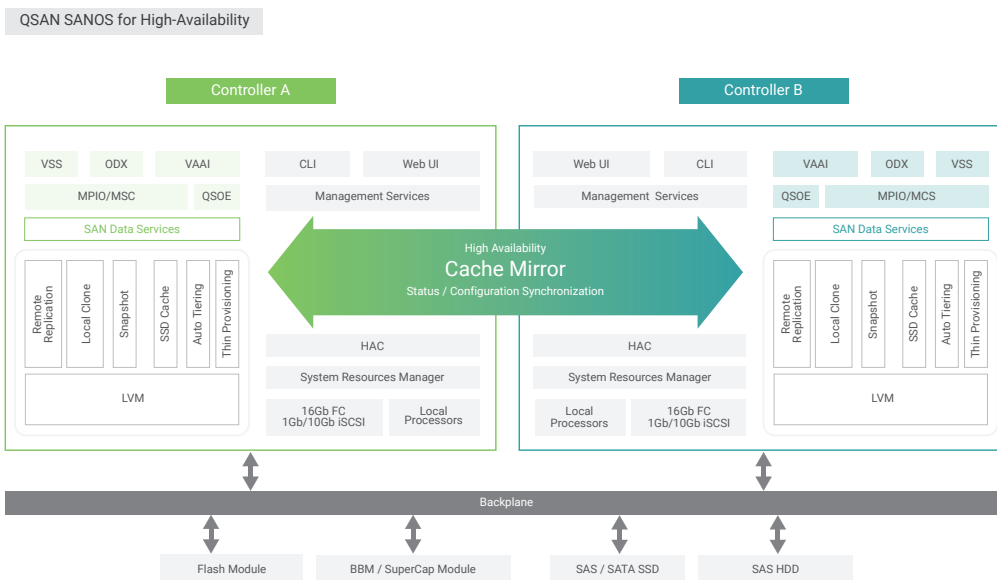


QSAN XS5226 model (2.5" 2U 26-bay) is the industry's first high-density all flash/low power SAN storage. It has accelerated performance, consolidated space, and two more storage bays (up to 8% more capacity) than the popular 2U 24-bay products available on the market. This can further lower the cost per TB, minimize IT rack space, and generate a smaller and greener footprint.

SANOS 4.0 SAN Operating System

SANOS 4.0 is QSAN's proprietary SAN storage operating system. SANOS 4.0 is equipped with a refreshingly simple to use web GUI and is easily deployable into any infrastructure.

Based on the Linux kernel, SANOS delivers comprehensive storage functionality including advanced storage management, complete RAID level protection, fast RAID rebuild, storage pool migration, thin provisioning, SSD cache, auto-tiering, snapshot, data backup and disaster recovery, virtualization support, performance monitoring, scale-up support, and more.



SANOS System Architecture

Advanced Storage Management

QSAN's in-house developed RAID stack technology has a proven track record of being deployed in enterprise environments for over ten years and is highly trusted by both SMB and enterprise customers globally. The advanced LVM (Logical Volume Manager) builds a solid foundation of disk virtualization to provide complete RAID level protection, enhanced performance, and many other enterprise-level storage features built on top.

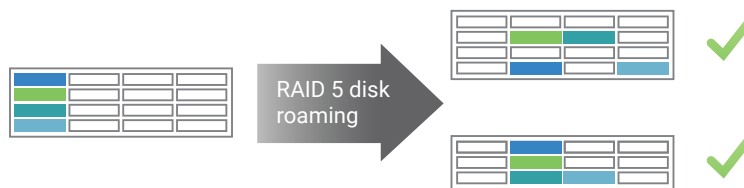
• Advanced RAID Level Technology

The most efficient and economical data protection method for enterprise and SMBs remains to be RAID technology. Using RAID has two advantages – high availability and better performance. SANOS supports complete RAID levels including RAID 0, 1, 0+1, 3, 5, 6, 10, 30, 50, 60, N-way mirror, and also RAID EE technology including RAID 5EE, 6EE, 50EE, 60EE for reducing the long rebuild risk. You can choose the appropriate RAID level to best suit the application requirements.

	RAID 0	RAID 1	RAID 3	RAID 5	RAID 6	N-way Mirror
Min.# Drives	1	2	3	3	4	3
Fault Tolerance	No protection	One drive failure	One drive failure	One drive failure	Two drive failures	N-1 drive failures
Read Performance	Very Good	Very Good	Very Good	Very Good	Very Good	Very Good
Write Performance	Excellent	Good	Good	Good	Fair to Good	Fair
Capacity Utilization (Min. - 26 drives)	100%	50%	67% ~ 96%	67% ~ 96%	50% ~ 92%	4% ~ 33%
	RAID 0+1	RAID 10	RAID 30	RAID 50	RAID 60	
Min.# Drives	4	4	6	6	8	
Fault Tolerance	One drive failure in each sub-array	One drive failure in each sub-array	One drive failure in each sub-array	One drive failure in each sub-array	Two drive failures in each sub-array	
Read Performance	Excellent	Excellent	Very Good	Very Good	Very Good	
Write Performance	Very Good	Very Good	Good	Good	Fair to Good	
Capacity Utilization (Min. - 26 drives)	50%	8% ~ 50%	67% ~ 92%	67% ~ 92%	50% ~ 85%	

• Intelligent Disk Roaming

The XS5200 SAN storage system will automatically recognize and locate member disks of a pool among enclosures. If you set a RAID group to offline in order to relocate the disk drives to another enclosure, there is no need to insert the disk drives at the exactly same slot as it was in the previous enclosure. Therefore, you don't need to purposely make room for the roaming disks on the target enclosure by moving disks, you can add them anywhere slots are vacant.



Intelligent Disk Roaming

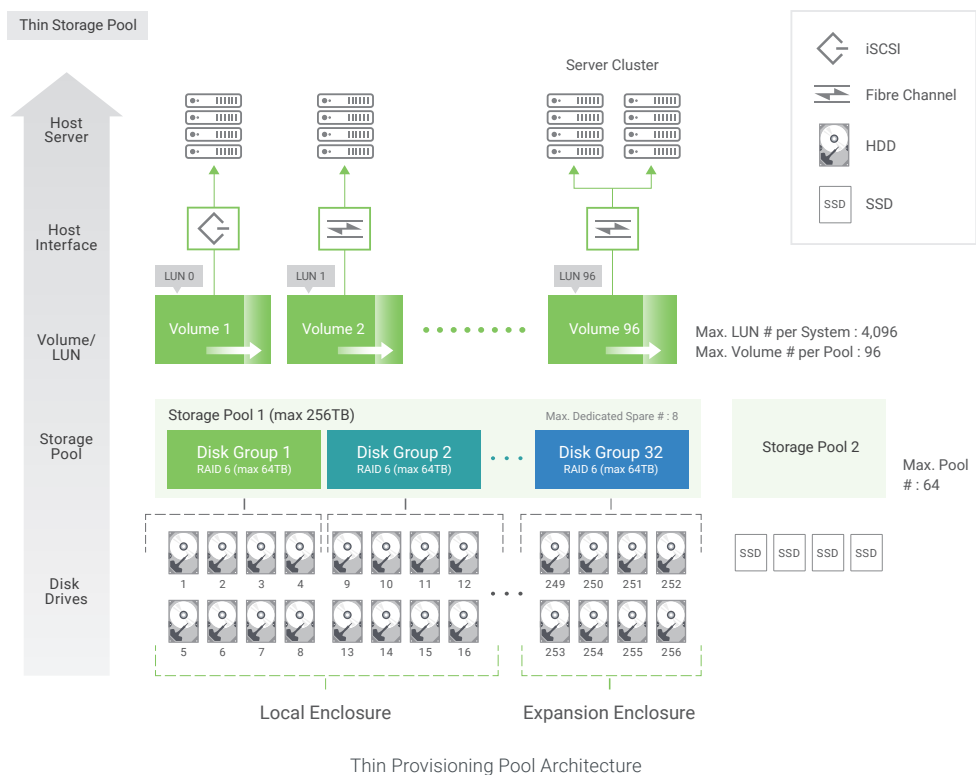
Fast RAID Rebuild Technology

As disk capacities of 8TB and larger become more commonplace, RAID rebuilding times become longer. The industry standard is 1TB per hour, so a full-driver or RAID group rebuild can take hours or even days to complete. QSAN fast RAID rebuild technology analyzes the volume structure by isolating data blocks and free blocks, so in the event of a failure the storage system will only rebuild the area of the RAID array in use. The Fast RAID rebuild feature also includes metadata maintenance operations to intelligently detect blocks that are no longer in use. This practice can reduce RAID rebuilding time by 50% or more.



SANOS 4.0 Storage Pool Architecture

QSAN storage pool supports a variety of SFF/LFF SAS disk drives and SFF SAS/SATA³ SSD drives. In thin provisioning, several disk drives are combined together to form a "Disk Group" with RAID protection. Then several disk groups can be combined to form a storage pool. A volume (virtual disk) is then created out of the storage pool and served to application servers over either iSCSI or Fibre Channel connections.

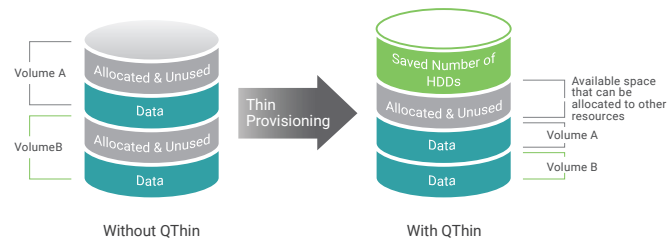


Storage Pool Type	Max Pool # per System	Max Disk Group # per Pool	Max Disk Drives # per Disk Group	Max Disk Drives # per Pool	Max Capacity per Disk Group	Max Capacity per Pool	Max Volume # per Pool	Max Volume Capacity	Max Volume # per System	Max LUN # per System	Max Host # per Controller for iSCSI	Max Host # per Controller for FC
Thin Provisioning Pool	64	32	64	256	64TB	256TB	96	256TB	4,096	4,096	512	256

³ 6G MUX board needed for 2.5" SATA drives in dual controller system.

Thin Provisioning (QThin)

Thin provisioning is a method of optimizing the efficiency with which the available space is utilized in SAN networks. In computing, thin provisioning involves using virtualization technology to give the appearance of having more physical resources than are actually available. Thin provisioning (QThin) operates by allocating disk storage space in a flexible manner among multiple users, based on the minimum space required by each user at any given time.



QThin uses our advanced storage pool architecture to achieve this functionality. The capacity of the storage pool is shared by all volumes inside the pool. This means that the XS5200 SAN is able to support more host servers with a fewer number of hard drives, thereby reducing initial acquisition cost. As the storage pool fills, you can add extra disk groups to expand the storage pool capacity without downtime.

• Policy-based Automatic Space Reclamation

Normally when data is deleted from the host or server, the unused storage blocks cannot be released back to the overall storage pool. QThin uses zero reclamation technology to recycle unused storage blocks to provide greater efficiency and a higher utilization rate. When enabled, the space reclamation process will run in the background automatically with the lowest system priority without affecting system performance.

You can set as many as six policies for each storage pool to define what to do when the predefined storage utilization threshold hits the limits. Starting with 60% utilization, the administrator can choose the level of warning and what action they would like to take, such as deleting snapshots, space reclamation, and de-activate the storage pool, allowing more flexibility and simpler administration.

Name	Status	Health	Total	Free	Available	Thin Provisioning	Auto Tiering	Volumes
QThin	Online	Good	1.45 TB	1.45 TB	1.45 TB	Enabled	Not Supported	0

No	Threshold	Level	Action
1	60%	Information	Take no Action
2	70%	Information	Take no Action Reclaim Space Delete Snapshots De-activate Pool
3	80%	Information	Reclaim Space Delete Snapshots De-activate Pool
4	85%	Warning	Reclaim Space
5	90%	Warning	Delete Snapshots
6	95%	Warning	De-activate Pool

Thin Provisioning Pool Policy

• Lower Initial Investment in Drives

With QThin technology, companies don't need to buy a large number of disks to manage data growth confidently. Start with the minimal number of disks, and then bring additional drives online as capacity demand grows.

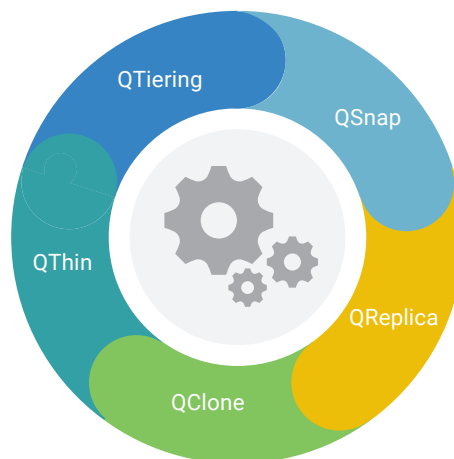


• Enhanced Storage Efficiency

QThin can release a volume's previously provisioned but unused capacity, helping to avoid a poor resource utilization rate and, in fact, drive utilization towards 100% with very little effort from system administrators. Available disk space can be made available to other hosts and servers. The XS5200 SAN system can serve more hosts and servers to achieve higher consolidation ratio. Therefore, QThin can significantly help reduce your total cost of ownership.

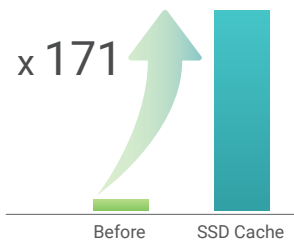
• Online Storage Pool Expansion and Fully-integrated with QSnap, QReplica, QCache, and QTiering

When QThin is enabled, the capacity of the storage pool can be expanded online by adding more disk groups with zero downtime. A thin storage pool can have up to 32 disk groups with each disk group containing up to 64 disk drives but up to 256 disk drives in a pool. QThin can also seamlessly integrate and interact with other QSAN storage features such as snapshot, local clone, remote replication, SSD cache, and auto tiering.



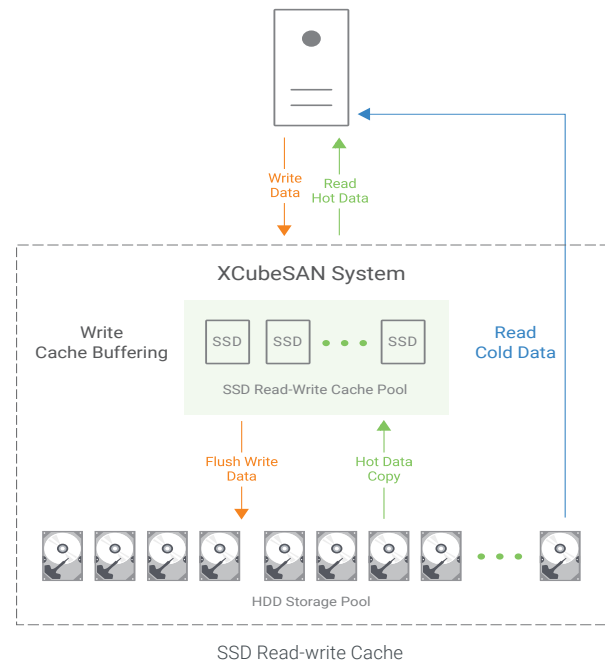
SSD Cache (QCache 2.0)

SSD cache is a large-capacity secondary cache that uses enterprise SSD drives positioned between the RAID controller's primary DRAM memory cache and hard disk drives (HDD). SSD read/write cache boosts random IOPS of read and write I/O for the system by copying frequently accessed random data to SSD drives, which are faster than HDDs. Therefore increasing the overall random IOPS and lowering the costs by using only few SSD drives. With this technology, QCache 2.0 can improve random read performance by up to 92 times and random write by up to 171 times. SSD drives also provide a much larger, scalable cache than the memory. The usable capacity of QCache is in proportion to the size of the controller system memory.



System Memory per Controller	Maximum SSD Cache Capacity per System
4GB	X (Not Support) ⁴
8GB	2TB
16GB	4TB
32GB	8TB
64GB	16TB

QCache 2.0 supports read cache and write cache⁵ which are up to four SSD cache pools per system. Each SSD cache pool can be used by one dedicated storage pool and its multiple volumes shared for effective usage of resources.



• Cache I/O Types

There are three defined cache I/O types and one customization option available which are applied to an SSD cache pool. According to your application, suitable cache I/O type would benefit the SSD running.

• Improve Performance, Spend Less

Statistically, only a portion of data is accessed frequently in any given storage tier and requires the higher performance of an SSD. This ratio of a small amount of SSD drives balanced with many HDDs offers the best performance (price/IOPS) at the lowest cost (price/GB) with optimal power efficiency (IOPS/kWh).

⁴ Please be aware that it needs 8GB to enable QCache.

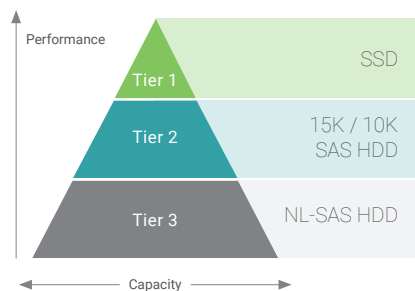
Auto Tiering (QTiering)

QSAN Auto Tiering (QTiering) cost-effectively and dynamically places hot data on SSD or faster hard drives and cold data on lower cost high-capacity drives, allowing you to optimize application performance without straining your budget or sacrificing capacity.

Our QTiering algorithm uses intelligent data analysis that continuously monitors data usage and ranks this data based on how often it is accessed. It will then use this information and make a decision on where your data should be.

Our intuitive SANOS 4.0 web UI interactively shows the data being gathered; how this data is being used, and how much of each tier storage should be assigned based on this information. Then at the scheduled time, the most accessed blocks that have been marked as "hot" data will be migrated into the highest performing tier, the least accessed or "cold" data will be migrated into the lowest cost - highest capacity drive tier.

All of this is managed in the background without user intervention. This tiered pool will also function the same as any standard QSAN pool, and access to our enterprise features such as QSnap and QReplica remains unchanged. This intelligent movement of data will allow the highest performance for the data you use the most, while keeping the total cost of ownership low and taking the burden of data management away from the IT organization.



3 Levels of Tiered Storage

QTiering supports 3 tiers of different drive types.

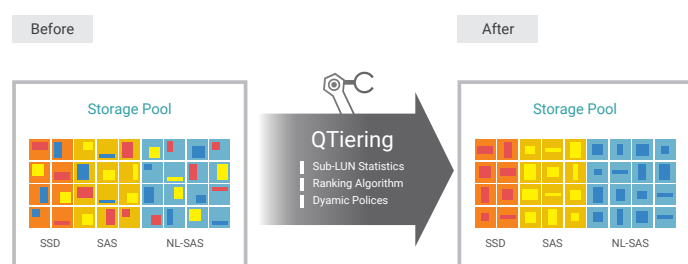
- Tier 1: SAS/SATA⁵ 2.5" SSD drives
- Tier 2: 15K/10K 3.5"/2.5" SAS drives
- Tier 3: 7.2K near-line 3.5"/2.5" SAS drives

QTiering manages the data relocation and monitors the data hotness ratio using half-life coefficient and advanced ranking mathematics. QTiering operates on three major functions.

- **Sub-LUN Statistics** : The volume is divided into 1GB units, which is called a Sub-LUN. This is the basic unit of data movement among tiers. Whenever there are I/O requests, the activity level of a Sub-LUN is determined by counting the read and write frequency to the Sub-LUN.
- **Ranking Algorithm** : Access records of each sub-LUN are collected and analyzed every hour. LVM maintains a cumulative I/O count and weighs each I/O by how recently it arrived, using a half-life coefficient. The ranking algorithm then uses these statistics to calculate the percentage of hot data.
- **Data Movement** : The data relocation engine then uses these percentages as guidance to move sub-LUNs between storage tiers automatically. The data relocation process will neither interfere with I/O nor stop I/O services. When data relocation begins to move sub-LUNs from slower tier to the faster tier, you will notice the I/O performance increases over time.

Dynamic Tiering Policies

There are 5 tiering policies available that are applied to a volume. They can be changed on the fly without affecting I/O services. Tiering policies not only affect the behavior of data relocation, but they also determine what the initial tier of the volume should be. These 5 tiering policies provide more flexibility and options to satisfy all deployment scenarios.



Intelligent QTiering Mechanism

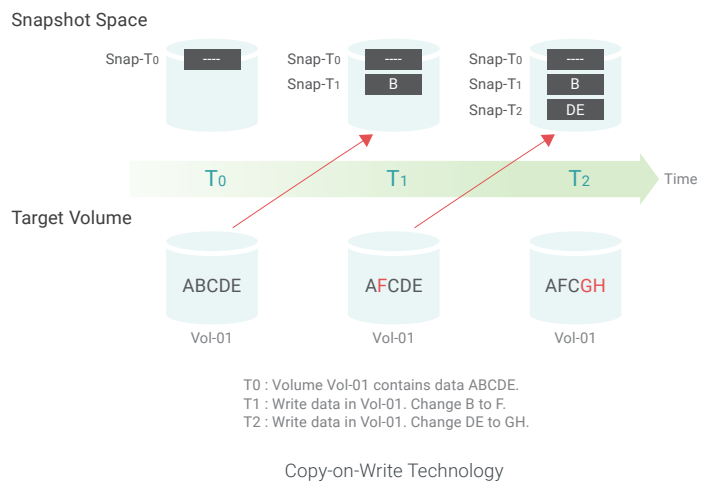
■ Most Activity ■ Normal Activity ■ Least Activity

⁵ 6G MUX board needed for 2.5" SATA drives in dual controller system.

Snapshot (QSnap)

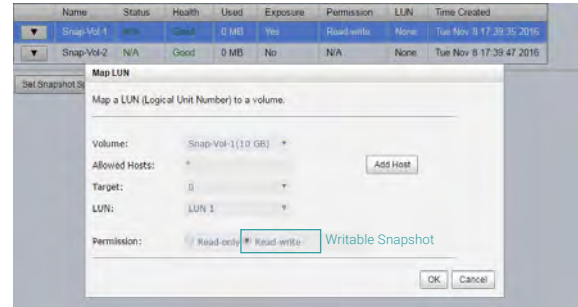
QSAN snapshot (QSnap) is based on copy-on-write technology. It's a block-based and differential backup mechanism. QSnap's functionality is designed to be highly efficient; it keeps a point in time record of block-level, incremental data changes of the target volume. QSnap can help recover a volume to a previous state quickly to meet enterprise SLA (Service Level Agreement) requirements of RPO (Recovery Point Objectives) and RTO (Recovery Time Objective).

QSnap is the easiest and most effective measurement to protect against ransomware attacks, virus attacks, accidental file deletion, accidental file modification, or unstable system hardware caused by bad I/O cable connection, unstable power supply, etc.



• Writable Snapshot Support

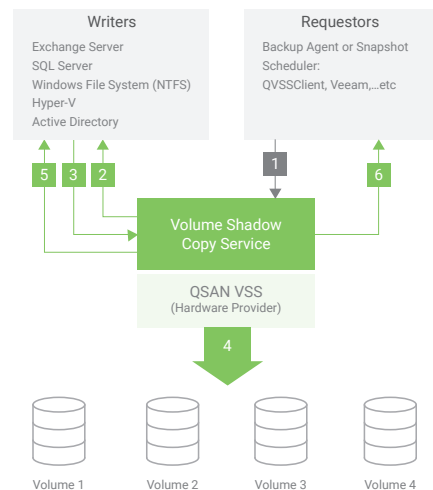
Apart from the rollback function, QSnap allows direct access to the snapshot content with read or read/write permissions. There are two benefits. One is that it will not consume the free capacity of the storage pool. The other one is that it will not affect the content of the target volume. Before mapping a LUN to the snapshot, the snapshot needs to be exposed to be prepared for accessing. An example of these benefits would be that programmers or developers can easily test a previous version of their compiled codes simply by mounting an older snapshot version onto a LUN instead of rolling back the snapshot and overwriting the existing source codes.



Writable Snapshot Support

• Integration with Windows VSS

QSnap is compatible with Windows VSS (Volume Shadow Copy Service). VSS is a host memory flush mechanism for creating consistent "point in time" copies of data known as "shadow copies". A Windows agent utility is provided to bridge and synchronize the information between the XS5200 SAN system and Windows operating system. After implementation, you can trigger a snapshot directly from Windows without any data consistency issues.



Volume Shadow Copy Service Workflow

High Reliability

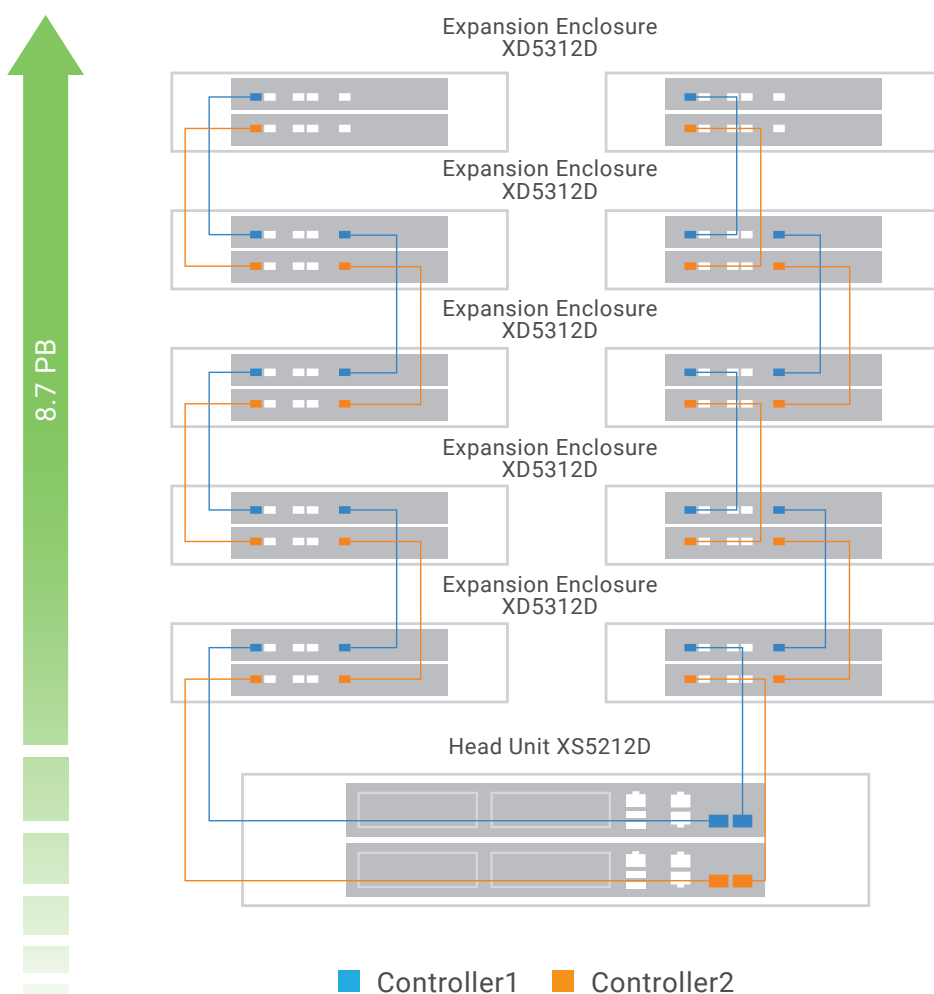
The XS5200 series uses the most reliable components from world-class manufacturers such as Intel® processors, QLogic Fibre Channel chips, branded power supply units, and DRAM modules to ensure the system has the highest level of quality and reliability.

In the product design and development phase, QSAN engineers have carried out a very thorough design review including circuit board's signal quality measurement and critical path analysis. QSAN engineers also perform a rigorous system level design verification, including the thermal cycling test, thermal shock test, shock & vibration test, drop test, humidity test, and EMI+EMS test.

Before mass production, the XCubeSAN XS5200 series has passed rigid factory-rigid Reliability Demonstration Testing, thermal cycle test, high-low temperature test inside the chamber and aging room to achieve over 150,000 hours of MTBF rating. The XS5200 series is truly a highly reliable and trust-worthy solution for enterprise and SMB markets.

Flexible Scale-up Solution

The XS5200 SAN provides a massive scale-up capability by connecting either our XD5300 expansion enclosures. It can support up to an impressive 626 disk drives or up to 8.7PB of raw storage space when using 15TB SAS drives.



Huge Scale-up Capacity Potential

· High Density, High Flexibility, and High Scalability

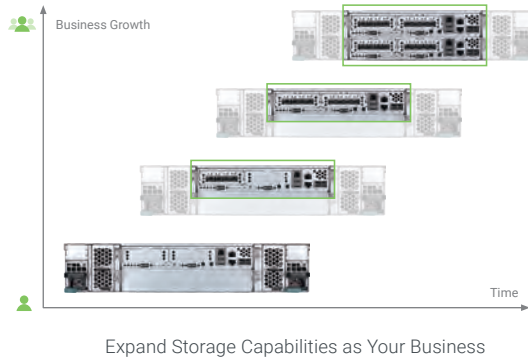
The expansion XD5300 series is the best scale-up solution for XCubeSAN XS5200 customers. It is available in 4U LFF 24-bay (XD5324), 3U 16-bay (XD5316), 2U 12-bay (XD5312), and the world's first 2U SFF 26-bay (XD5326). These can support up to ten expansion enclosures. For ultimate flexibility, there is no restriction on mixing different form factors of expansion enclosures together. Users can choose different form factors of expansion enclosures in accordance with their needs and budget.



SAN Models	Expansion Enclosures	Max. No. of Expansion Units	Max. No. of Disk Drives	Max. Raw Capacity	
				(LFF 10TB, SFF 2.4TB)	(LFF 14TB, SFF 15.36TB)
XS5212 (2U 12-bay)	XD5312 (2U 12-bay)	10	12 + 12 x 10 = 132	1,320TB	1,848TB
	XD5316 (3U 16-bay)	10	12 + 16 x 10 = 172	1,720TB	1,408TB
	XD5324 (4U 24-bay)	10	12 + 24 x 10 = 252	2,520TB	3,528TB
	XD5326 (2U 26-bay)	10	12 + 26 x 10 = 272	744TB	4,161.6TB
	XD5360 (4U 60-bay)	10	12 + 60 x 10 = 612	6,120TB	8,568TB
XS5216 (3U 16-bay)	XD5312 (2U 12-bay)	10	16 + 12 x 10 = 136	1,360TB	1,904TB
	XD5316 (3U 16-bay)	10	16 + 16 x 10 = 176	1,760TB	2,464TB
	XD5324 (4U 24-bay)	10	16 + 24 x 10 = 256	2,560TB	3,584TB
	XD5326 (2U 26-bay)	10	16 + 26 x 10 = 276	784TB	4,217.6TB
	XD5360 (4U 60-bay)	10	16 + 60 x 10 = 616	6,160TB	8,624TB
XS5224 (4U 24-bay)	XD5312 (2U 12-bay)	10	24 + 12 x 10 = 144	1,440TB	2,016TB
	XD5316 (3U 16-bay)	10	24 + 16 x 10 = 184	1,840TB	2,576TB
	XD5324 (4U 24-bay)	10	24 + 24 x 10 = 264	2,640TB	3,696TB
	XD5326 (2U 26-bay)	10	24 + 26 x 10 = 284	864TB	4,329.6TB
	XD5360 (4U 60-bay)	10	24 + 60 x 10 = 624	6,240TB	8,736TB
XS5226 (2U 26-bay)	XD5312 (2U 12-bay)	10	26 + 12 x 10 = 146	1,262.4TB	2,079.36TB
	XD5316 (3U 16-bay)	10	26 + 16 x 10 = 186	1,662.4TB	2,639.36TB
	XD5324 (4U 24-bay)	10	26 + 24 x 10 = 266	1,462.4TB	3,759.36TB
	XD5326 (2U 26-bay)	10	26 + 26 x 10 = 286	686.4TB	4,392.96TB
	XD5360 (4U 60-bay)	10	26 + 60 x 10 = 626	6,062.4TB	8,799.36TB

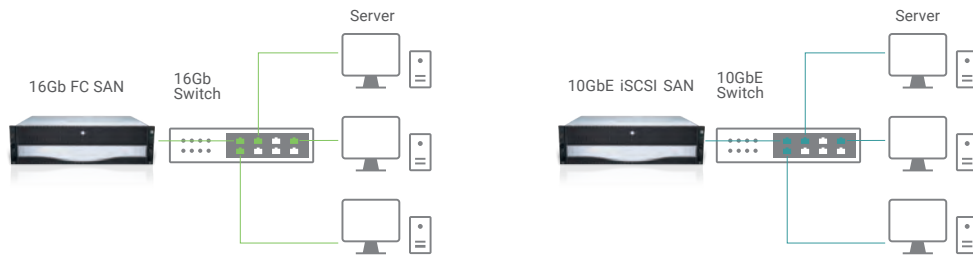
Modular I/O Ports for iSCSI SAN and Fibre Channel SAN

Each XS5200 SAN controller has two host card slots that can be configured as iSCSI SAN, Fibre Channel, or a mix of both. There are various types of optional host cards available to match your specific need, including 16Gb Fibre Channel, 10GbE iSCSI, and 1GbE iSCSI. You can choose the appropriate host card for your initial requirement, as your business grows you can add extra host cards to enhance system connectivity.

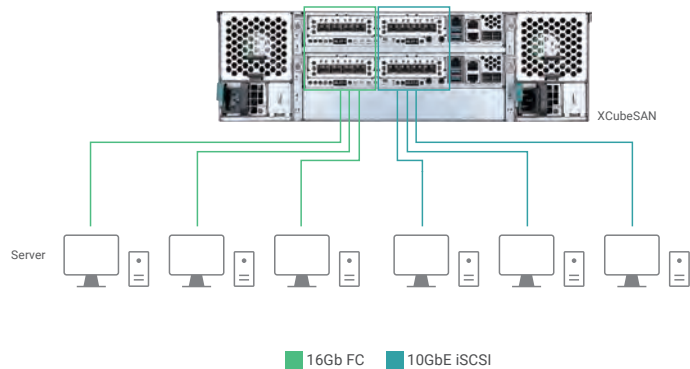


In dual controller configurations, the system can support up to 20 ports of 10GbE iSCSI or 8 ports of 16Gb FC connectivity⁶. Both 10GbE iSCSI with SFP+ interface and 10GBASE-T interface can coexist and deliver concurrent data services and maximize uptime. Utilizing these numerous ports, you can cost-effectively connect multiple host servers directly to the XS5200 series without using an FC switch or Ethernet switch.

Standard Deployment with Switches



XCubeSAN Deployment without Switches
More servers can direct attach to SAN Storage



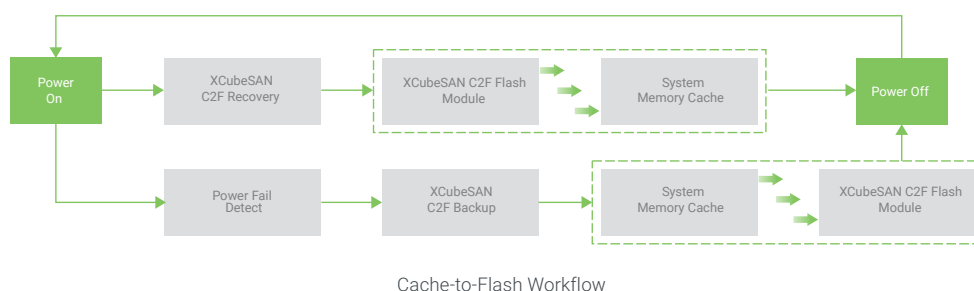
⁶ Slot 1 is where the optional either 4/2 x 16Gb FC, 4/2 x 10GbE iSCSI, or 4 x 1GbE iSCSI connection modules will install. In addition, slot 2 is where the optional either 4/2 x 16Gb FC (20Gb bandwidth), 4/2 x 10GbE iSCSI (20Gb bandwidth), or 4 x 1GbE iSCSI connection modules only.

Cache-to-Flash Memory Protection Technology

In the event of power loss, the I/O cache data stored in volatile memory will be lost, this can cause data inconsistencies, especially in database applications. The XS5200 series can provide an optional Cache-to-Flash memory protection function that will safely transfer the memory cache data to a non-volatile flash device for permanent preservation. The optional Cache-to-Flash module comes with an M.2 flash module and either a BBM (Battery Backup Module) or an SCM (Super Capacitor Module).

Reliable and Advanced Protection Mechanism

The following is the working sequence of QSAN Cache-to-Flash mechanism.




Cache-to-Flash technology will first flush CPU cache to memory RAM, then flush memory RAM to the M.2 flash module to maintain the upmost data consistency. It leverages the strength of both the BIOS and CPU to quickly backup memory RAM data to the flash module. In order to quickly move data from memory RAM to the flash module, M.2 PCI-Express interface flash module is selected for better performance and less power consumption. In Cache-to-Flash recovery phase, BIOS will check C2F flag status. If C2F flag is ON, I/O cache data will be recovered from the M.2 flash module and then continue normal booting. If C2F flag is OFF, the normal booting process continues. Compared to the traditional BBM solution of 72 hours standard, Cache-to-Flash technology is more efficient, less risky, and consumes less power.

Hot Pluggable Design with Zero Downtime

The optional Cache-to-Flash module comes with an M.2 flash module and either a BBM or an SCM. All modules are hot pluggable with zero system downtime for an extra measure of availability and reliability. The M.2 flash module can be plugged in on the left hand side from the rear of the chassis. The Power module can be plugged in on the right hand side from the rear of the chassis.

The BBM can protect all memory capacity. But the SCM can protect up to 16GB memory per controller. If your system memory is higher than 16GB, please select BBM solution.

Device	Protection Memory Capacity
Battery Backup Module + Flash Module	Protect all memory capacity
Super Capacitor Module + Flash Module	Protect up to 16GB memory per controller



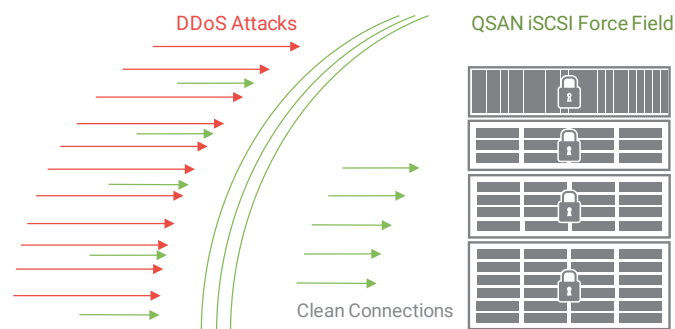
Cache-to-Flash Modules

Comprehensive Data Security

Data security and data integrity are one of the most important subjects for enterprises and SMBs. SANOS 4.0 provides the most comprehensive features to guarantee your data security and protect your businesses from data theft, unauthorized disclosure, malicious network attack, and accidental corruption.

· iSCSI Force Field against Mutant DDoS Attack

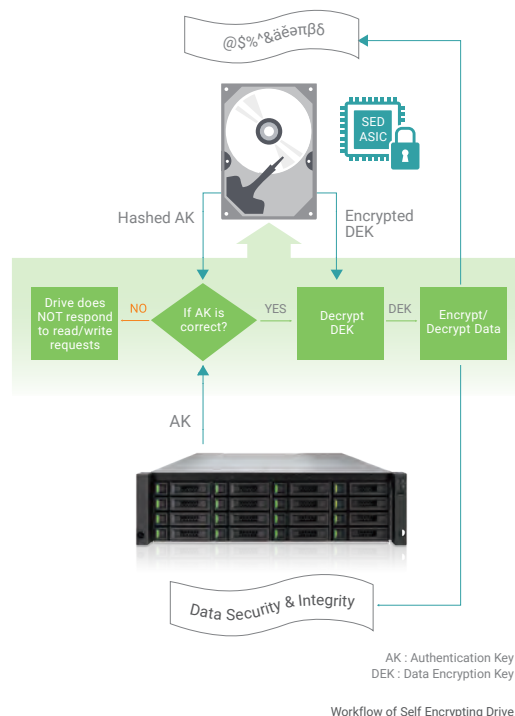
iSCSI connections over the Internet are prone to all sorts of network attacks. The XS5200 series has gone through rigorous network-attack tests using Mu Dynamics Mu-8000 appliance. DoS (Denial of Service) tests and mutation tests based on FUZZing technology are conducted to verify that the iSCSI target protocol stack is robust enough and smart enough to fend off all kinds of malicious attempts. It's a serious threat that demands serious solutions. QSAN iSCSI Force Field can guarantee your data has the highest level of security ensuring business continuity.



· SED (Self-Encrypting Drive) Support

SANOS supports Self-Encrypting Drives. SEDs have a circuit built into the disk drive controller chip that encrypts all data to the magnetic media and decrypts all the data from the media automatically.

SED is the perfect technology to stop the increasingly frequent loss or theft of sensitive data through careless disposal of unused or failed drives containing sensitive data. The XS5200 series can adopt SED drives for better data security without system overhead.



• iSCSI CHAP Authentication

The XS5200 series supports iSCSI CHAP authentication. iSCSI CHAP and mutual CHAP use a challenge-response mechanism to verify connection authentication and provide access control to prevent unrelated initiators (servers) from accessing storage resources.

• Login Lock & Auto Logout

To avoid concurrent modifications to the same storage resource, login lock function is provided to allow only a single IP address login at a time. The second login attempt will be blocked by login lock function.

Auto Logout function will automatically log out the current web GUI session after a preset idle time limit. This minimizes the risk of an unattended IT computer from unauthorized access.

Login Options	
Auto Logout:	Disabled
Login Lock:	Disabled 5 minutes 30 minutes 1 hour
Apply	

Login Security Setting

Data Backup & Disaster Recovery

SANOS 4.0 provides complete array-based backup functions including local backup and remote replication. They can build up a solid safety net for disaster recovery and meet the enterprise RPO and RTO requirement.

• Local Volume Clone (QClone)

Local volume clone (QClone) is used to make a duplicate copy of a volume in the same storage pool as well as in a separate storage pool within the same enclosure. In setting up a local clone task, the first clone is a full copy. From that point on, the cloning is a differential copy, created using QSAN's snapshot functionality. Manual and scheduled tasks are available for management flexibility. In the event that the source volume fails, IT managers can quickly switch to the cloned volume and resume data services.



Local Clone or Volume Copy

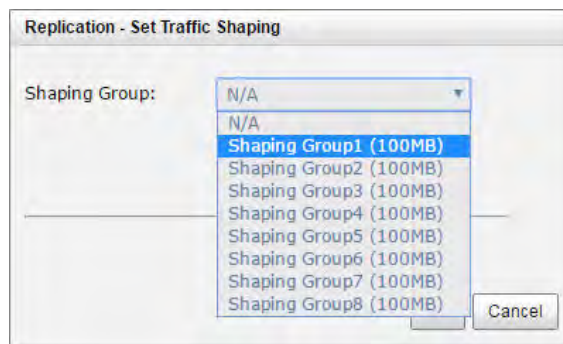
Remote Replication (QReplca 2.0)

QSAN remote replication is a block-level, asynchronous, differential remote volume backup function through LAN or WAN. QReplca 2.0 has many powerful capabilities including unlimited bandwidth, traffic shaping, and multiple connections per replication task. It's the most cost-effective and efficient way to perform remote data backup. All XCubeSAN storage systems running SANOS 4.0 or above can replicate to each other using QReplca 2.0 completely free of charge.



Remote Replication through WAN

QReplca 2.0 uses the iSCSI function to set up a replication connection. It can use the full bandwidth of the assigned network port to allow optimum backup speed. However, in order to balance replication traffic and non-replication traffic, the traffic shaping function can help to reserve necessary bandwidth for non-replication I/O.



Traffic Shaping in QReplca 2.0

If the replication task requires more bandwidth, QReplca 2.0 allows multiple connections per task by intelligently balancing the backup task across multiple connections to enhance the bandwidth.

Task:

No.	Source Volume	Status	%	Shaping	Speed	Target Volume	Capacity	Schedule	Time
1	Source-Vol-1	Online		N/A	0 MB	Target-Vol-2	11.00 GB	N/A	Tue

Task 'Source-Vol-1' Path:

No.	Source Port	Target IP Address	Target Name	LUN	Status
1	Auto	10.10.1.101	iqn.2004-08.com.qsan:dev0.ctr1	0	Connected
2	Auto	10.10.1.103	iqn.2004-08.com.qsan:dev0.ctr2	0	Connected

Create Rebuild Replication Options Traffic Shaping Configuration

Task 'Source-Vol-1' Path:

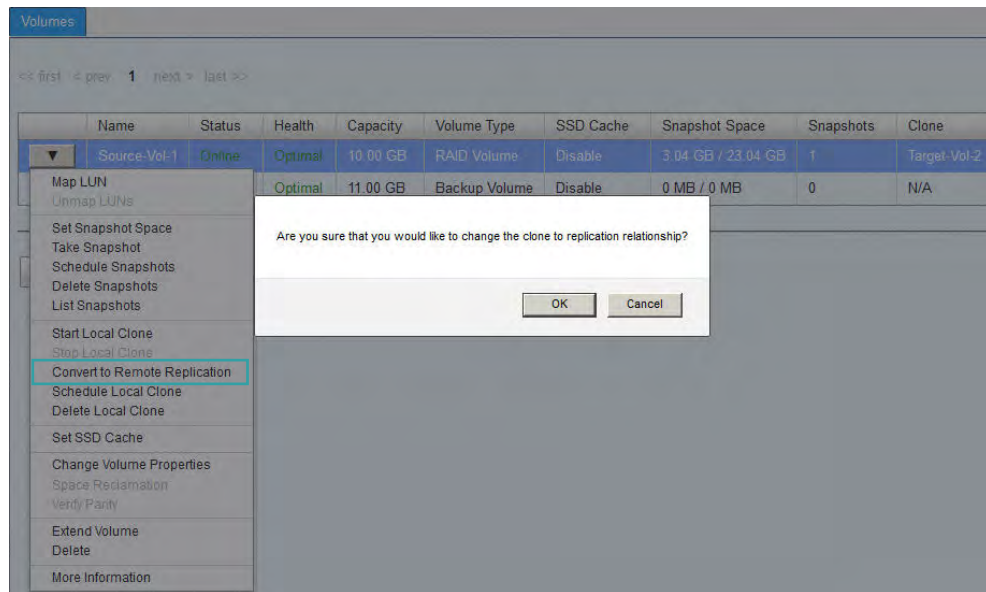
No.	Source Port	Target IP Address	Target Name	LUN	Status
1	Auto	10.10.1.101	iqn.2004-08.com.qsan:dev0.ctr1	0	Connected
	Auto	10.10.1.102			
2	Auto	10.10.1.103	iqn.2004-08.com.qsan:dev0.ctr2	0	Connected

Multiple Connections per Remote Replication Task

· QClone to QReplica

Both manual and scheduled replication tasks are supported for flexible management. To handle huge remote replication (for example 60TB) tasks, QReplica 2.0 allows transforming a local clone task into a remote replication task. You can perform the local volume clone first for the full copy, then use the disk roaming function to physically transport the disk drives that contain the cloned volume to the remote site. Lastly, use QReplica 2.0 to transform a local clone task to a remote replication one.

Remote backup for 60TB data from New York City to Boston Comparison of "QClone to QReplica" and "100Mbps Internet"		
Method	Total transmission time	Used bandwidth
100Mbps Internet	About 55 days	60TB
QClone to QReplica	1 day	0

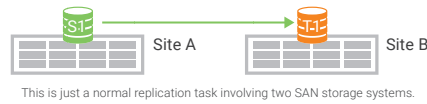


Convert QClone Task to a QReplica Task

Remote Replication Topologies

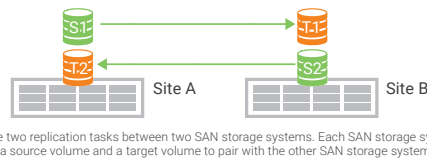
The XS5200 series can support multiple topologies to suit various disaster recovery configurations. They are one-directional, bi-directional, one-to-many, and many-to-one. Both the source volume and target volume in a replication connection are exclusive to the pair. Each SAN storage system in the XS5200 series can support up to 32 replication tasks concurrently. Below are the supported topologies.

One-Directional

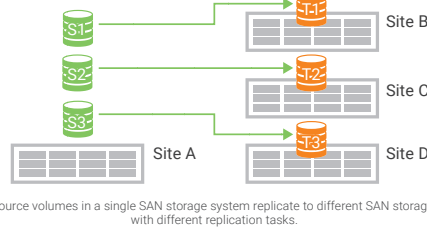


S1, S2, S3: Source Volumes
T1, T2, T3: Target Volumes

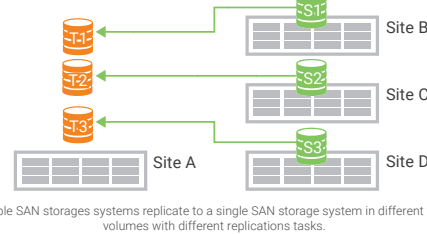
Bi-Directional



One-to-Many



Many-to-One



QReplica 2.0 Topologies

Online Volume Restoration

In the event of a problem with a RAID group after suffering multiple disk failures, QSAN volume restoration can help to salvage the volume to the extent possible by reverting storage pool configuration and volume configuration to the previous healthy one. Sometimes the volumes can be salvaged, and all the data can then be replicated to a different storage pool for safety. It is strongly recommended to replace the problematic disk drives after all the data is replicated. This feature is very unique to XCubeSAN, and has been proven in action to help businesses recover mission-critical data when all else failed.

Pool Name	RAID	Volume	Volume Capacity	Disks Used	Disk Slot	Time	Event Logs
123	RAID 0	123	100 GB	1	0:14	2015/03/19 10:45:36 CST	Disk is removed from the system.
123	RAID 0	123	100 GB	2	0:14, 0:15	2015/03/19 10:12:54 CST	The volume is created.
qq	RAID 0+1	qq	10 GB	4	0:14	2016/02/17 17:02:29 CST	The volume is created.
qq	RAID 0+1	q	446 GB	4	0:14	2016/02/17 15:31:18 CST	The volume is created.

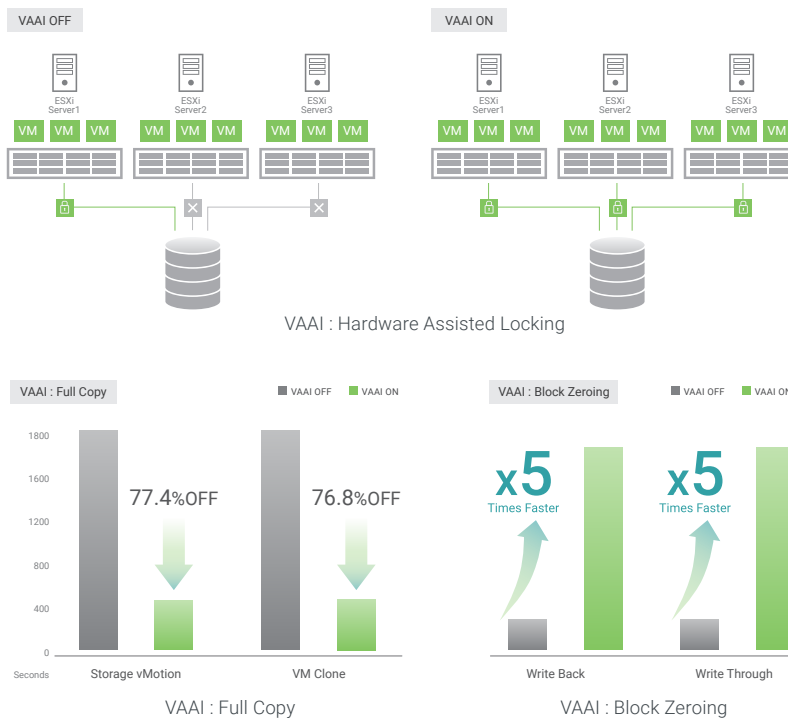
Volume Restoration via SANOS 4.0

Virtualization

The XS5200 series is virtualization-aware and is highly integrated with leading hypervisor platforms; It is certified by the latest VMware vSphere with VAAI support, Windows Server 2016/2012 R2 with Windows ODX support, and the latest Citrix XenServer. All of these features make the XS5200 series an ideal primary storage system for virtualized datacenters to help provision, migrate, and manage VM storage faster and more efficiently, allowing hypervisors to provision data intelligently and run even more VMs with a lower total cost of ownership.

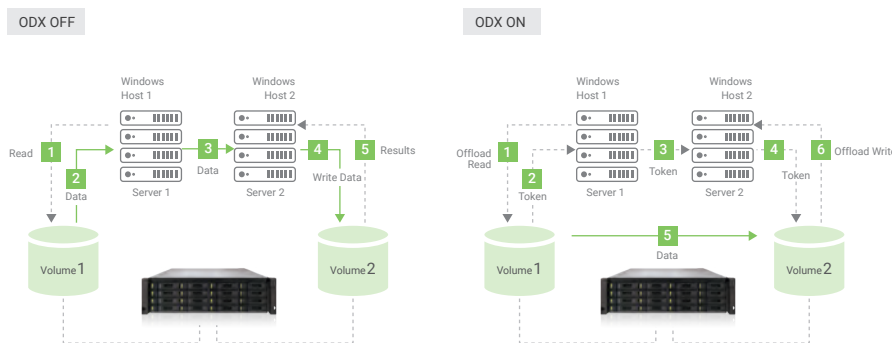
• SANOS 4.0 for VMware VAAI Support

SANOS 4.0 supports VMware VAAI functionality. VAAI is a set of APIs that allow ESXi hosts to offload specific virtual machine and storage management tasks to the SAN. With the support of SANOS 4.0 VAAI, it can offload the ESXi system overheads and as a result the ESXi system can reduce the usage of its CPU, memory, and storage fabric bandwidth. VAAI for iSCSI and FC supports Hardware Assisted Locking, Block Zero, Full Copy, and Thin Provisioning with space reclaim. Therefore, the XS5200 series can greatly enhance the performance of operations in a virtualized environment.



• SANOS 4.0 for Microsoft Hyper-V Support

With ODX (Offloaded Data Transfer) support in SANOS 4.0, the XS5200 series becomes a high performance iSCSI storage solution in Windows Server 2016/2012 R2 and Hyper-V virtualized environments. It highly reduces the loading of Windows host servers and improves the performance of copy and move operations.



Microsoft Windows ODX (Offloaded Data Transfer)

Interoperability

The XS5200 series has high interoperability, it supports host server running the following operating systems :

- Windows Server 2008, 2008 R2, 2012, 2012 R2, 2016
- SLES (SUSE Linux Enterprise Server) 10, 11, 12
- RHEL (Red Hat Enterprise Linux) 5, 6, 7
- CentOS (Community ENTERprise Operating System) 6, 7
- Solaris 10, 11
- FreeBSD 9, 10
- Mac OS X 10.11 or later

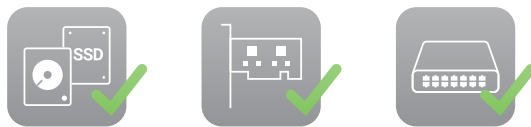


Wide Compatibility

QSAN's product development and test departments have invested extensive testing resources to verify compatibility with peripherals including SAS HDDs, SSDs, switches, HBA cards, and third-party backup utilities to make sure there is maximum compatibility with the XS5200 series and typical add-on devices.

In addition, QSAN has qualified the most popular SAS HDDs and SSDs on the market. Our users have the highest level of flexibility to choose from the most advanced HDDs or SSD flash drives and purchase the best cost/performance drives to suit their needs, significantly reducing IT infrastructure investment costs.

Detailed compatibility matrix can be found at QSAN website: <https://qsan.com/rdht4a>



Optional USB LCM

QSAN systems have an innovative way of simply displaying system information. Our portable USB LCM (LCD Module) can help you identify management port IP address, facilitate basic setting, and shutdown the system. When you have finished the task at hand, simply unplug the USB LCM and then you are free to plug it into the next XCubeSAN system to configure or monitor that system. This can prevent unauthorized attempts to find management port IP address or shutdown the system intentionally. This feature means that the XCubeSAN is secure from internal tampering as well as being easy to configure. USB LCM is an optional item.

Green Technology

At QSAN, we pride ourselves on our commitment to build highly efficient and low carbon footprint devices. To safeguard the earth and our environment, the XS5200 product range uses various green technologies for energy savings and minimization of your carbon footprint.



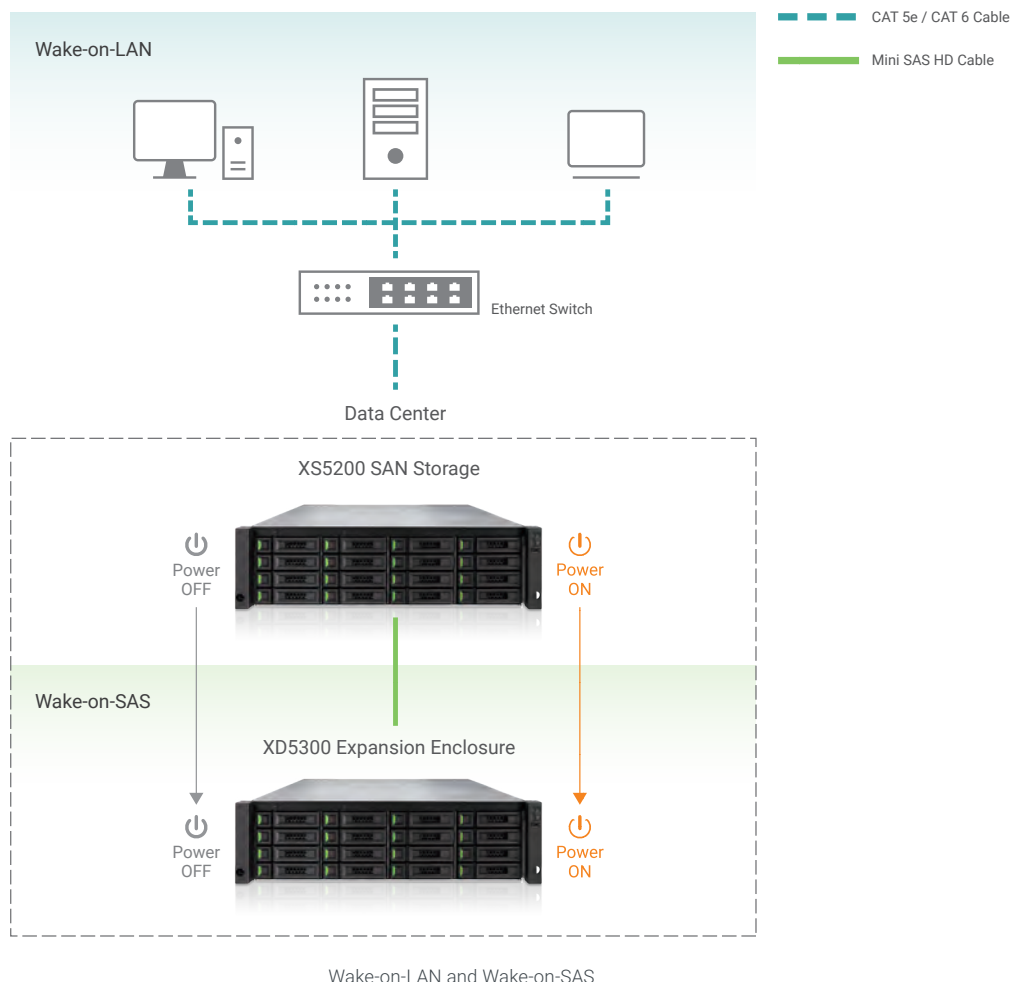
• 80 PLUS Platinum Power Supply

The XS5200 product family are all equipped with dual redundant 80 PLUS Platinum power supply modules for the ultimate in energy efficiency. At 50% load, these power supplies can provide up to 92% efficiency converting power from AC to DC. This efficiency means that our power supplies can greatly reduce the system's power loss and heat generation.

• Wake-on-SAS Technology

QSAN's Wake-on-SAS technology allows you to remotely power on/off all cascaded XD5300 expansion enclosures by using QSAN proprietary SAS cables. You can power on the XS5200 SAN system remotely using the Wake-on-LAN feature. The XS5200 can work with any available Wake-on-LAN freeware and shareware. Wake-on-SAS ensures that expansion enclosures will not run idly, consuming electricity after the SAN storage system is shut down for maintenance or other purposes. Wake-on-SAS can avoid unnecessary electricity waste by allowing your devices to be on only when it is necessary. A further advantage of Wake-on-SAS is that when you turn on the SAN, the expansion enclosures will wake automatically, so there is no need to worry about degrading a volume if you forget to turn them on first.

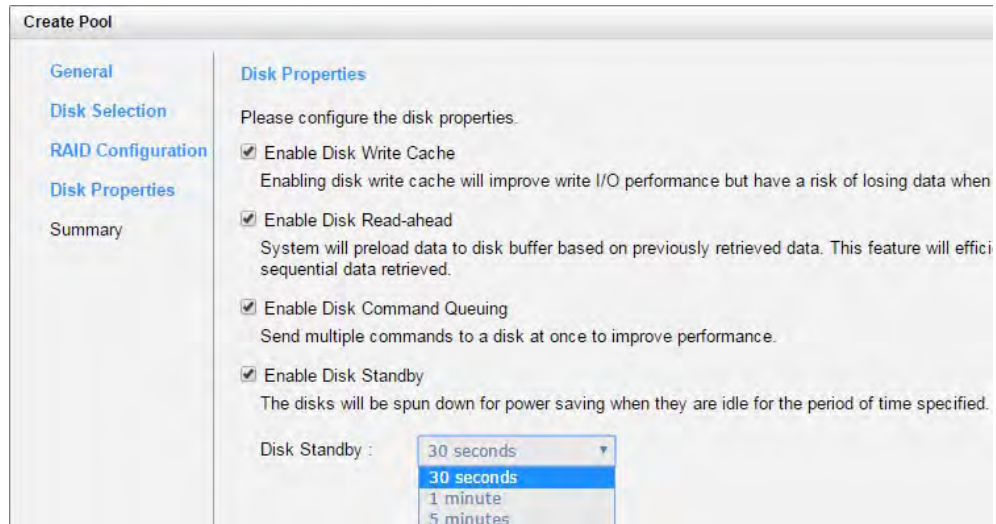
The following diagram shows that after the XS5200 SAN storage receives the magic packet being sent from the client's computer, it powers on automatically and uses Wake-on-SAS functionality to power-on the attached XD5300 expansion enclosures.



• Auto Disk Spin-down

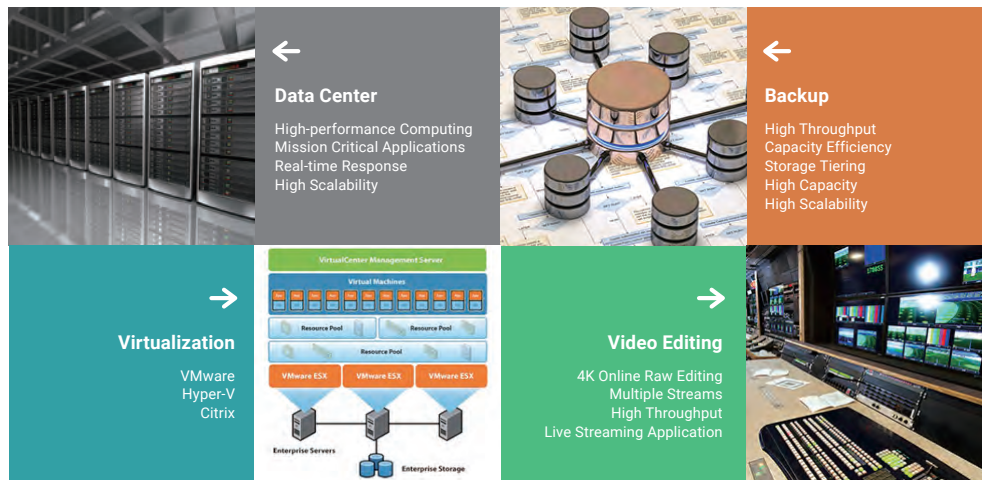
When the disk drives have not been accessed or are not scheduled to be accessed for a set time defined by you, the auto disk spin down functionality will kick in, and will power down the whole storage pool into a hibernation state, this gives you the freedom to optimize the energy usage of the XS5200 SAN storage system. Disk spin-down can save 60% to 70% power in low use scenarios such as backups.

When the system detects read/write requests coming in, it will issue a wake command, this takes a few seconds to spin up all the disk drives in the storage pool to resume all data services. The performance penalty can be off-set with SSD cache, however.



Auto Disk Spin-down Setup Menu

XCubeSAN Applications



· Enable Mission Critical Applications

Mission critical applications such as OLTP (On-Line Transaction Processing) applications, virtualized data center and data warehouses, HPC (High Performance Computing), and e-mail, web, file serving in large enterprise usually require direct access to drives (block-level devices), heavy I/O workloads and low I/O latency. With QSOE 2.0 (QSAN Storage Optimization Engine) technology and next generation storage platform, the XS5200 series provides extreme IOPS and is accelerated performance enterprise-class storage systems designed for these I/O critical applications. Leveraging the powerful Intel® Xeon® multi-core processor to offload iSCSI stack and Fibre Channel stack to a dedicated core to accelerate performance and reduce I/O latency. QSOE 2.0 can boost performance by 50% and deliver stable high performance to enable all your mission critical business applications and remove your performance bottleneck.

· Live Database Management

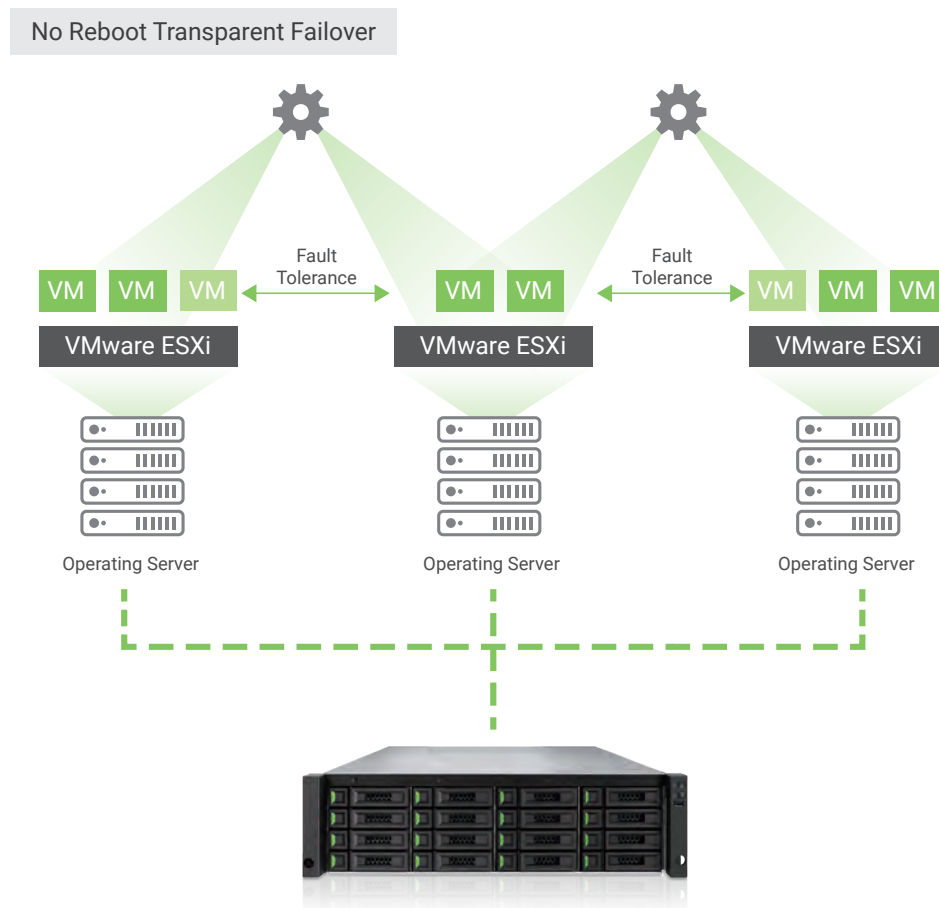
The XS5200 series is built for mission-critical database applications. Its five 9's high availability, dual active controllers, Cache-to-Flash memory protection, and efficient array-based backup solutions are all tailor-made for database applications. By using QCache or all flash solutions, database latency can increase performance up to 400%. Thanks to this performance increase, you can consolidate more databases into the XS5200 series and reduce database licensing fees by up to 50%. Using the free QSnap/QClone/QReplica enterprise storage functions reduces database backup time from hours to minutes by utilizing incremental copy-on-write technology. Running your SQL Server environment or Oracle 12c database environment on the XS5200 series means:

- Accelerated database performance with stunning throughput and high random IOPS
- Reduce or eliminate planned/unplanned downtime utilizing the redundant components, high availability and non-disruptive operations
- Reduce storage and data management costs and increase ROI (Return On Investment)
- Accelerate application test/develop process and deployment



· Server Virtualization Solution

The XS5200 series is a virtualization-ready SAN storage system. It supports vMotion, DRS, HA from VMware and Live Migration from Hyper-V and XenMotion from Citrix. You need a Dual-Active SAN storage solution to construct a fail-safe server virtualization deployment. Because all the VM image files are stored in the XS5200 SAN storage protected by the Dual-Active high availability. No matter if it is the failure of VM or a physical server or storage system, the server virtualization solution has no single point-of-failure from head to toe and can provide non-stop services all the time. The XS5200 Dual-Active SAN storage is most definitely “best practice” for server virtualization.



XS5200 dual-active SAN storage is the best practice for server virtualization




























· Video Editing

The business community at large has embraced HD video. Digital video is rapidly evolving to 4K and beyond, but adoption and consumption of ultra HD video creates immense stress and challenges for networks and storage on the whole. Next generation storage networking such as 16Gb Fibre Channel and 10GbE iSCSI are ideal tools to solve 2K/4K video entertainment requirements such as content ingest, editing, rendering, streaming and broadcasting. The XS5200 series offers super wide bandwidth up to 8 x 16Gb FC ports and 20 x 10GbE iSCSI ports through innovative dual host card design. Stable 12,000MB/s throughput and over 400K random write IOPS can support multiple concurrent streams of FullHD 1080, UltraHD 2K, 4K, and even 3D video content without dropping a single frame.



The XS5200 series is an ideal storage solution for pre-production and post-production environments. Superior performance and high throughput make the XS5200 series work seamlessly with popular video editing applications such as Adobe Creative Suite, AVID Media Composer, Apple Final Cut Pro, etc.

Accessories /

Model Name	Picture	Description	Applied Models	
DIM-D44GB		DDR4 ECC 4GB Memory Module		
DIM-D48GB		DDR4 ECC 8GB Memory Module		
DIM-D416GB		DDR4 ECC 16GB Memory Module		
DIM-D432GB		DDR4 ECC 32GB Memory Module		
C2F-BM128G		Cache-to-Flash Module: C2F-BBMD + C2F-FLHMD		
C2F-SP128G		Cache-to-Flash Module: C2F-SP16G + C2F-FLHMD (Protect up to 16GB memory per controller)		
C2F-BBMD		Battery Backup Module for Cache-to-Flash		
C2F-SP16G		Super Capacitor Module for Cache-to-Flash (Protect up to 16GB memory per controller)		
C2F-FLHMD		Flash Module for Cache-to-Flash		
HQ-16F4S2 HQ-16F2S2		4-port 16Gb Fibre Channel Host Card (SFP+) 2-port 16Gb Fibre Channel Host Card (SFP+)		XS5224
HQ-10G4S2 HQ-10G2T		4-port 10GbE iSCSI Host Card (SFP+) 2-port 10GBASE-T iSCSI Host Card (RJ45)	XS5216 XS5212	
HQ-01G4T		4-port 1GBASE-T iSCSI Host Card (RJ45)	XS5226	
GBC-SFP+16Gb-J		16G Fibre Channel SFP+ Optical Transceiver		
GBC-SFP+10Gb-F		10GBASE-SR SFP+ Optical Transceiver		
GBC-SFP+8Gb-F		8G Fibre Channel SFP+ Optical Transceiver		
CBL-OPL500		Optical FC Cable, LC-LC, 5 Meters		
CBL-OPL200		Optical FC Cable, LC-LC, 2 Meters		
CBL-CNL		Console Cable, Phone-jack, 1 Meter		
CBL-UPS		UPS Cable, Phone-jack, 1 Meter		
CBL-12SW150		SAS 12G Expansion Cable with Wake-on-SAS, SFF-8644 to SFF-8644, 1.5 Meters (This cable is not suitable for connecting HBA or RAID card)		
CBL-12SH150		SAS 12G Expansion Cable, SFF-8644 to SFF-8644, 1.5 Meters		
LCM-U162		Portable USB LCM		
SLR-RM3640		Slide Rail		
HDT-351		3.5" Disk Drive Tray	XS5224	
HDM-351		SATA 6Gb MUX Board and Bracket for HDT-351 (2.5" SATA drives only)	XS5216 XS5212	
HDT-251		2.5" Disk Drive Tray	XS5226	
HDM-251		SATA 6Gb MUX Board and Bracket for HDT-251		

Hardware Specifications /



Model Name	XS5224D (Dual)	XS5216D (Dual)	XS5212D (Dual) XS5212S (Single)	XS5226D (Dual) XS5226S (Single)
Form Factor	4U 24-bay, LFF	3U 16-bay, LFF	2U 12-bay, LFF	2U 26-bay, SFF
RAID Controller	Dual-active or Single-upgradable controller			
Processor	Intel® Xeon® 4-core processor (8-core processor models are also available)			
Memory (per Controller)	DDR4 ECC 8GB, up to 128GB (four DIMM slots, insert two DIMMs or more will boost performance)			
Host Connectivity (per Controller)	Host Card Slot 1 (optional) : 4 x 16Gb FC (SFP+) ports 2 x 16Gb FC (SFP+) ports 4 x 10GbE iSCSI (SFP+) ports 2 x 10GbE iSCSI (RJ45) ports 4 x 1GbE iSCSI (RJ45) ports		Host Card Slot 2 (optional) : 4 x 16Gb FC (SFP+) ports ⁷ 2 x 16Gb FC (SFP+) ports ⁷ 4 x 10GbE iSCSI (SFP+) ports ⁷ 2 x 10GbE iSCSI (RJ45) ports 4 x 1GbE iSCSI (RJ45) ports	
	Built-in 2 x 10GBASE-T iSCSI (RJ45) ports Built-in 1 x 1GbE management port			
Expansion Connectivity (per Controller)	Built-in 2 x 12Gb/s SAS wide ports (SFF-8644)			
Drive Type	Mix & match 3.5" & 2.5" SAS, NL-SAS, SED HDD 2.5" SAS, SATA ⁸ SSD			2.5" SAS, NL-SAS, SED HDD 2.5" SAS, SATA ⁸ SSD
Expansion Capabilities	Up to 10 expansion units using XD5300 series 12Gb SAS expansion enclosure XD5324 (LFF 24-bay), XD5316 (LFF 16-bay), XD5312 (LFF 12-bay), XD5326 (SFF 26-bay)			
Max. Drives Supported	624	616	612	626
Dimension (H x W x D)	19" Rackmount 170.3 x 438 x 515 mm	19" Rackmount 130.4 x 438 x 515 mm	19" Rackmount 88 x 438 x 515 mm	19" Rackmount 88 x 438 x 491 mm
Memory Protection	Cache-to-Flash module (optional) Battery backup module + Flash module (To protect all memory capacity) Super capacitor module + Flash module (To protect up to 16GB memory per controller)			
LCM (LCD Module)	USB LCM (optional)			
Power Supply	770W/850W 1+1 redundant 80 PLUS Platinum			
Fan Module	2 x hot pluggable/redundant fan modules			

Warranty

Warranty	System : 3 years Battery backup module : 1 year Super capacitor module : 1 year
Regulatory	CE, FCC, BSMI, VCCI, KCC

Operating Environment

Temperature	Operating temperature : 0 to 40°C Shipping temperature : -10°C to 50°C
Relative Humidity	Operating relative humidity : 20% to 80% non-condensing Non-operating relative humidity : 10% to 90%

⁷ Slot 2 provides 20Gb bandwidth.

⁸ 6Gb MUX board needed for 2.5" SATA drives in dual controller system.

Software /

Operating System

- 64bit embedded Linux

Storage Management

- RAID level 0, 1, 0+1, 3, 5, 6, 10, 30, 50, 60, and N-way mirror
- RAID EE level 5EE, 6EE, 50EE, and 60EE
- Flexible storage pool ownership
- Thin Provisioning (QThin) with space reclamation
- SSD Cache (QCache⁹)
- Auto Tiering (QTiering⁹)
- Global, local, and dedicated hot spares
- Write-through and write-back cache policy
- Online disk roaming
- Spreading RAID disk drives across enclosures
- Background I/O priority setting
- Instant RAID volume availability
- Fast RAID rebuild
- Online storage pool expansion
- Online volume extension
- Online volume migration¹⁰
- Auto volume rebuilding
- Instant volume restoration
- Online RAID level migration
- SED & ISE drive support
- Video editing mode for enhanced performance
- Disk drive health check and S.M.A.R.T attributes
- Storage pool parity check and media scan for disk scrubbing
- SSD wear lifetime indicator
- Disk drive firmware batch update

iSCSI Host Connectivity

- Proven QSOE 2.0 optimization engine
- CHAP & mutual CHAP authentication
- SCSI-3 PR (Persistent Reservation for I/O fencing) support
- iSNS support
- VLAN (Virtual LAN) support
- Jumbo frame (9,000 bytes) support
- Up to 256 iSCSI targets
- Up to 512 hosts per controller
- Up to 1,024 sessions per controller

Fibre Channel Host Connectivity

- Proven QSOE 2.0 optimization engine
- FCP-2 & FCP-3 support
- Auto detect link speed and topology
- Topology supports point-to-point¹¹ and loop
- Up to 256 hosts per controller

High Availability

- Dual-Active (Active/Active) SAN controllers
- Cache mirroring through NTB bus
- ALUA support
- Management port seamless failover
- Fault-tolerant and redundant modular components for SAN controller, PSU, FAN module, and dual port disk drive interface
- Dual-ported HDD tray connector
- Multipath I/O and load balancing support (MPIO, MC/S, Trunking, and LACP)
- Firmware update with zero system downtime

Security

- Secured Web (HTTPS), SSH (Secure Shell)
- iSCSI Force Field to protect from mutant network attack
- iSCSI CHAP & mutual CHAP authentication
- SED & ISE drive support

Storage Efficiency

- Thin Provisioning (QThin) with space reclamation
- Auto Tiering (QTiering⁹) with 3 levels of storage tiers

Networking

- DHCP, Static IP, NTP, Trunking, LACP, VLAN, Jumbo frame (up to 9,000 bytes)

Advanced Data Protection

- Snapshot (QSnap), block-level, differential backup
- Writeable snapshot support
- Manual or schedule tasks
- Up to 64 snapshots per volume
- Up to 64 volumes for snapshot
- Up to 4,096 snapshots per system
- Remote Replication (QReplica)
- Asynchronous, block-level, differential backup based on snapshot technology
- Traffic shaping for dynamic bandwidth controller
- Manual or schedule tasks
- Auto rollback to previous version if current replication fails
- Up to 32 schedule tasks per controller
- Volume clone for local replication
- Configurable N-way mirroring
- Integration with Windows VSS (Volume Shadow Copy Service)
- Instant volume restoration
- Cache-to-Flash memory protection⁹
- M.2 flash module
- Power module: BBM or SCM (Super Capacitor Module)
- USB and network UPS support with SNMP management

Virtualization Certification

- Server Virtualization & Clustering
- Latest VMware vSphere certification
- VMware VAAI for iSCSI & FC
- Windows Server 2016, 2012 R2 Hyper-V certification
- Microsoft ODX
- Latest Citrix XenServer certification

Easy Management

- USB LCM⁹, serial console support, online firmware update
- Intuitive Web management UI, secured web (HTTPS), SSH (Secured Shell), LED indicators
- S.E.S. support, S.M.A.R.T. support, Wake-on-LAN, and Wake-on-SAS

Green & Energy Efficiency

- 80 PLUS Platinum power supply
- Wake-on-LAN to turn on or wake up the system only when necessary
- Auto disk spin-down

Host Operating Systems Support

- Windows Server 2008, 2008 R2, 2012, 2012 R2, 2016
- SLES 10, 11, 12
- RHEL 5, 6, 7
- CentOS 6, 7
- Solaris 10, 11
- FreeBSD 9, 10
- Mac OS X 10.11 or later

⁹ The function is optional and is not included in the default package.

¹⁰ The feature is based on RAID level migration of disk groups on the fly in thick provisioning pools.

¹¹ 16Gb Fibre Channel only supports Point-to-Point topology.



QSAN Technology, Inc. | Learn more by visiting www.qsan.com

Address : 4F., No.103, RuiHu Street, NeiHu District, Taipei, Taiwan 114 Email : sales@qsan.com Telephone : +886-2-7720-2118 Fax : +886-2-7720-0295